

# Professional Apache Tomcat

Amit Bakore  
Debashish Bhattacharjee  
Sandip Bhattacharya  
Vivek Chopra  
Chad Fowler  
Ben Galbraith  
Romin Irani  
Sing Li  
Chanoch Wiggers





# Summary of Contents

Chapter 1:	<b>Apache and Jakarta Tomcat</b>	1
Chapter 2:	<b>JSP and Servlets</b>	15
Chapter 3:	<b>Tomcat Installation</b>	31
Chapter 4:	<b>Tomcat Installation Directory and Architecture</b>	65
Chapter 5:	<b>Basic Tomcat Configuration</b>	77
Chapter 6:	<b>Web Application Administration</b>	109
Chapter 7:	<b>Manager Configuration</b>	131
Chapter 8:	<b>Advanced Tomcat Features</b>	163
Chapter 9:	<b>Class Loaders</b>	201
Chapter 10:	<b>HTTP Connectors</b>	221
Chapter 11:	<b>Web Server Connectors</b>	235
Chapter 12:	<b>The WARP Connector</b>	243
Chapter 13:	<b>The AJP Connector</b>	257
Chapter 14:	<b>Tomcat and IIS</b>	285
Chapter 15:	<b>JDBC Connectivity</b>	311
Chapter 16:	<b>Tomcat Security</b>	335
Chapter 17:	<b>Additional Uses for Ant</b>	381
Chapter 18:	<b>Log4J</b>	405
Chapter 19:	<b>Shared Tomcat Hosting</b>	425
Chapter 20:	<b>Server Load Testing</b>	461
Appendix A:	<b>Axis</b>	491
Appendix B:	<b>Apache SSL Setup</b>	497
Index		513

# Professional Apache Tomcat

Amit Bakore  
Debashish Bhattacharjee  
Sandip Bhattacharya  
Vivek Chopra  
Chad Fowler  
Ben Galbraith  
Romin Irani  
Sing Li  
Chanoch Wiggers



# Professional Apache Tomcat

Published by  
**Wiley Publishing, Inc.**  
10475 Crosspoint Boulevard  
Indianapolis, IN 46256  
www.wiley.com

Copyright © 2003 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

Library of Congress Card Number: 2003107064

ISBN: 0-7645-4372-5

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

1B/QV/QW/QT/IN

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8700. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4447, E-Mail: permcoordinator@wiley.com.

**LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY:** WHILE THE PUBLISHER AND AUTHOR HAVE USED THEIR BEST EFFORTS IN PREPARING THIS BOOK, THEY MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS BOOK AND SPECIFICALLY DISCLAIM ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES OR WRITTEN SALES MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A PROFESSIONAL WHERE APPROPRIATE. NEITHER THE PUBLISHER NOR AUTHOR SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

For general information on our other products and services or to obtain technical support, please contact our Customer Care Department within the U.S. at (800) 762-2974, outside the U.S. at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

**Trademarks:** Wiley, the Wiley Publishing logo, Wrox, the Wrox logo, the Wrox Programmer to Programmer logo and related trade dress are trademarks or registered trademarks of Wiley in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

# Trademark Acknowledgements

Wrox has endeavored to provide trademark information about all the companies and products mentioned in this book by the appropriate use of capitals. However, Wrox cannot guarantee the accuracy of this information.

## Credits

### Authors

Amit Bakore  
Debashish Bhattacharjee  
Sandip Bhattacharjee  
Vivek Chopra  
Chad Fowler  
Ben Galbraith  
Romin Irani  
Sing Li  
Chanoch Wiggers

### Additional Material

Matthew Moodie  
Peter Wainwright

### Technical Reviewers

Subrahmanyam  
Allamaraja  
Kapil Apshankar  
Steve Baker  
Yogesh Bhandarkar  
Vivek Chopra  
Kris Dahl  
Richard Huss  
Romin Irani  
Meeraj Kunnumpurath  
Massimo Nardone  
Neil Matthew  
Richard Stones  
Sameer Tyagi  
Keith Wannamaker

### Commissioning Editors

Paul Cooper  
Ranjeet Wadhvani

### Technical Editors

Kedar Kamat  
Unnati Kulkarni  
Matthew Moodie  
Nilesh Parmer

### Author Agents

Shivanand Nadkarni  
Safiulla Shakir

### Indexers

Bill Johncocks

### Proof Reader

Agnes Wiggers

### Production Coordinators

Rachel Taylor  
Pip Wonson

### Illustrations

Santosh Haware  
Manjiri Karande

### Cover

Dawn Chellingworth

# About the Authors

## Amit Bakore

Amit is a Sun Certified Java Programmer with a couple of other certifications. Currently based in Pune (India), he works as a System Analyst for Sansui Software. He has been working mainly on J2EE and XML on Linux. Before landing in the world of software, he graduated from Pune University, with Electronics as a major, with first class.

*I humbly dedicate this work to 'HIM' and his parents, Dr.Ramkrishna & Sau. Vaijayanti. I sincerely thank all my friends, colleagues, and wellwishers for their extensive support and guidance.*

## Debashish Bhattacharjee

Debashish Bhattacharjee is a Principal Consultant with the Management Consulting Services unit of PriceWaterhouseCoopers. He has 10 years of experience implementing projects for Fortune 500 clients in the United States and Canada. His areas of expertise are systems integration and project management. He has served as chief architect and led technical teams tasked with the implementation of e-commerce applications, portal implementations, web infrastructure, ERP, and client-server applications.

In his role as consultant, Debashish is often responsible for advising clients on best practices and the adoption of technology. He is the published author of several industry articles.

## Sandip Bhattacharjee

Sandip is an open source enthusiast and an active participant in various open source communities in India, especially his local LUG – Indian Linux Users Group, Delhi (ILUGD), <http://www.linux-delhi.org/>. He has been programming right from his school days in 1991, and carries an engineering degree in Textile technology and an MBA in marketing.

He has been professionally involved in open source technologies for the past three years. He is currently a freelance programmer and advises businesses on ways to use the open source revolution to their advantage.

## Vivek Chopra

Vivek has eight years of experience in software design and development, the last two years of which have been in web services and various XML technologies. He is the co-author of *Professional ebXML Foundations* (ISBN 186100-590-3) and *Professional XML Web Services* (ISBN 1-86100-509-1) both from Wrox Press. He is also a committer for UDDI4J, an open source Java API for UDDI. His other areas of experience and interest include compilers, middleware, clustering, GNU/Linux, and mobile computing. He is currently consulting in the domain area of web services.

Vivek holds a Bachelor's degree in Electronics and a Master's in Computer Science, both from Pune University, India. He lives and works in the beautiful San Francisco Bay Area, and can be reached at [vivek@soaprpc.com](mailto:vivek@soaprpc.com).

## Chad Fowler

Chad Fowler is CTO of GE Appliances Bangalore, India office. For the past four years, he has been an active advocate of open source Java technologies in the enterprise, revolving around the Enhydra suite of software. Driven into software development by a less-than-healthy addiction to the video game Doom, he dropped his professional music career and never looked back. His current interests focus on the Ruby programming language, learning the Hindi (spoken/written – not programming) language, and Agile Software Development methodologies.

## Ben Galbraith

Before graduating from High School, Ben Galbraith was hired by a major Silicon Valley computer manufacturer to develop Windows-based client-server applications with international deployments and hundreds of users. In 1995, he began developing for the Web and fell in love with Unix, vi, and Perl. After building countless web applications with Perl, Ben discovered server-side Java in 1999 and his relationship with Perl has since become somewhat estranged.

He is presently a consultant in Provo, Utah. He regularly lectures, evangelizes, and gives classes on Java technology. Ben has no college degree but if he had the time he would study both ancient and modern history.

## Romin Irani

Romin Irani is a Senior Software Engineer with InSync Information Systems, Inc in Fremont, California. He graduated with a Bachelors degree in Computer Engineering from University of Bombay, India. He has around seven years of experience, starting out in the Microsoft world but now fully immersed into Java technologies. He welcomes your comments at [romin@rocketmail.com](mailto:romin@rocketmail.com).

*I am most thankful to my wife Devayani, whose cooperation and love made this possible. And of course due credits to my parents, Khushru and Gulrukh for all that they have taught me in life.*

## Sing Li

First bitten by the computer bug in 1978, Sing has grown up with the microprocessor and the Internet revolution. His first PC was a \$99 do-it-yourself COSMIC ELF computer with 256 bytes of memory and a 1 bit LED display. For two decades, Sing has been an active author, consultant, instructor, entrepreneur, and speaker. His wide-ranging experience spans distributed architectures, web services, multi-tiered server systems, computer telephony, universal messaging, and embedded systems.

Sing has been credited with writing the very first article on the Internet Global Phone, delivering voice over IP long before it became a common reality. Sing has participated in several Wrox projects in the past, has been working with (and writing about) Java, Jini, and JXTA since their very first available releases, and is an active evangelist for the unlimited potential of P2P technology.



## **Chanoch Wiggers**

Chanoch is a Java Programmer working with JSP and Servlets who until recently worked at Wrox Press as a Technical Architect (writing this stuff is even more fun than writing about this stuff). Chanoch would like to thank the reviewers and the guys at Wrox Press, especially Shivanand Nadkarni.





# Table of Contents

<b>Chapter 1: Apache and Jakarta Tomcat</b>	<b>1</b>
<b>Humble Beginnings: The Apache Project</b>	<b>2</b>
<b>The Apache Software Foundation</b>	<b>3</b>
Apache Projects	3
The Jakarta Project	4
<b>Distributing Tomcat</b>	<b>6</b>
<b>The Big Picture: J2EE</b>	<b>8</b>
<b>Using Tomcat with the Apache Web Server</b>	<b>11</b>
<b>Summary</b>	<b>13</b>
<b>Chapter 2: JSP and Servlets</b>	<b>15</b>
<b>First Came CGI...</b>	<b>15</b>
<b>Then Servlets Were Born...</b>	<b>16</b>
<b>And On To JSPs...</b>	<b>20</b>
<b>Web Application Architecture</b>	<b>26</b>
<b>Java Site Architecture</b>	<b>27</b>
<b>Summary</b>	<b>29</b>
<b>Chapter 3: Tomcat Installation</b>	<b>31</b>
<b>Installing a JVM</b>	<b>31</b>
<b>Tweaking the JVM for Performance</b>	<b>36</b>
<b>Tomcat Installation</b>	<b>38</b>
Tomcat Windows Installer	39
Installing Tomcat On Windows Using the ZIP File	47
Installing Tomcat from Source	48
Installing Tomcat On Linux	48
Running Tomcat with the Server Option	51

## Table of Contents

---

<b>The Tomcat Installation Directory</b>	<b>52</b>
<b>Ant Installation</b>	<b>54</b>
Installing Ant On Windows	54
Installing Ant On Linux	55
<b>Installing Tomcat from Source</b>	<b>55</b>
The Ant Build File	57
Building Tomcat	59
<b>Troubleshooting and Tips</b>	<b>60</b>
<b>Summary</b>	<b>63</b>
<b>Chapter 4: Tomcat Installation Directory and Architecture</b>	<b>65</b>
<b>The Installation Directory</b>	<b>65</b>
The bin Directory	66
The classes, lib, server, and common Directories	67
The conf Directory	67
The logs Directory	67
The webapps Directory	68
<b>Tomcat Architecture</b>	<b>69</b>
The Server	70
The Service	71
The Connectors	71
The Engine	72
The Remaining Classes in the Tomcat Architecture	74
Summary	74
<b>Chapter 5: Basic Tomcat Configuration</b>	<b>77</b>
<b>Component-Based Configuration</b>	<b>78</b>
<b>Files in \$CATALINA_HOME/conf</b>	<b>78</b>
<b>Basic Server Configuration</b>	<b>79</b>
Authentication and the tomcat-users.xml File	97
The Default Deployment Descriptor – conf/web.xml	97
Fine-Grained Access Control: catalina.policy	103
<b>Summary</b>	<b>107</b>
<b>Chapter 6: Web Application Administration</b>	<b>109</b>
<b>The Contents of a Web Application</b>	<b>109</b>
URL Mappings	110
Public Resources	111
The META-INF Folder	112
The WEB-INF Folder	113
The classes Folder	113
The tlds Folder	113
The lib Folder	114

---

<b>The web.xml File</b>	<b>114</b>
The XML Header	116
The DTD Declaration	116
<web-app>	116
<b>Summary</b>	<b>128</b>
<hr/> <b>Chapter 7: Manager Configuration</b>	<hr/> <b>131</b>
<b>Sample Web Application</b>	<b>131</b>
<b>Tomcat 3.x Administration Tool</b>	<b>132</b>
Enabling Permissions for the Admin Tool	133
admin Application Tasks	133
<b>Tomcat 4.x Manager Application</b>	<b>137</b>
Enabling Access To the Manager Application	139
Manager Application Configuration	140
Manager Application Commands	142
<b>Tomcat Web Application Manager (4.1.7 Beta Only)</b>	<b>150</b>
Managing Applications with Ant (Tomcat 4.1 Only)	154
<b>Tomcat Administration Tool (Tomcat 4.1 Only)</b>	<b>158</b>
Admin Application Configuration	159
<b>The Future</b>	<b>160</b>
<b>Summary</b>	<b>160</b>
<hr/> <b>Chapter 8: Advanced Tomcat Features</b>	<hr/> <b>163</b>
<b>Valves – Interception Tomcat Style</b>	<b>164</b>
<b>Standard Valves</b>	<b>164</b>
<b>Access Logs Implementation</b>	<b>165</b>
<b>Single Sign-On Implementation</b>	<b>169</b>
<b>Restricting Access Via a Request Filter</b>	<b>172</b>
<b>Persistent Sessions</b>	<b>175</b>
Configuring a Persistent Session Manager	176
<b>JNDI Resource Configuration</b>	<b>178</b>
<b>Realms</b>	<b>188</b>
What Is a Realm?	188
Container Managed Security	189
Configuring JDBC Realms	190
Configuring JNDI Realms	193
Memory Realm	194
UserDatabase as a Realm	196
<b>Summary</b>	<b>198</b>

<b>Chapter 9: Class Loaders</b>	<b>201</b>
<b>Class Loader Overview</b>	<b>202</b>
Standard J2SE Class Loaders	203
Class Loader Attributes	207
Creating a Custom Class Loader	207
<b>Security and Class Loaders</b>	<b>208</b>
Class Loader Delegation	209
Core Class Restriction	209
Separate Class Loader Namespaces	209
Security Manager	209
<b>Tomcat and Class Loaders</b>	<b>210</b>
System Class Loader	211
Common Class Loader	211
Catalina Class Loader	213
Shared Class Loader	214
Web Application Class Loader	215
<b>Dynamic Class Reloading</b>	<b>215</b>
<b>Common Class Loader Pitfalls</b>	<b>216</b>
<b>Summary</b>	<b>218</b>
<b>Chapter 10: HTTP Connectors</b>	<b>221</b>
<b>HTTP Connectors</b>	<b>222</b>
Tomcat 3.x: HTTP/1.0 Connector	222
Tomcat 4.0: HTTP/1.1 Connector	225
Tomcat 4.1: Coyote HTTP/1.1 Connector	226
<b>Running Tomcat Behind a Proxy Server</b>	<b>230</b>
<b>Using Coyote HTTP with Tomcat 3.3.x</b>	<b>231</b>
<b>Using Coyote HTTP with Tomcat 4.0</b>	<b>231</b>
<b>Performance Tuning</b>	<b>232</b>
<b>Summary</b>	<b>233</b>
<b>Chapter 11: Web Server Connectors</b>	<b>235</b>
<b>Connector Architecture</b>	<b>236</b>
Connector Protocols	237
WARP Protocol	237
AJP Protocol	238
<b>Choosing a Connector</b>	<b>239</b>
JServ	239
jk	240
jk2	240
webapp	240
<b>Summary</b>	<b>241</b>

---

<b>Chapter 12: The WARP Connector</b>	<b>243</b>
<b>Introducing webapp</b>	<b>243</b>
<b>webapp Configuration</b>	<b>244</b>
webapp Binaries	244
Building webapp from Source	245
Configuration Changes in Apache	249
Configuration Changes in Tomcat	251
Testing the Installation	253
<b>webapp Bugs and Issues</b>	<b>253</b>
<b>Summary</b>	<b>255</b>
<b>Chapter 13: The AJP Connector</b>	<b>257</b>
<b>mod_jk</b>	<b>257</b>
The Apache JServ Protocol	258
The AJP Connector	258
Worker Implementations with mod_jk	259
Plug-In vs. In-Process	259
Multiple Tomcat Workers	260
Getting mod_jk	260
<b>Integrating Tomcat with Apache</b>	<b>262</b>
Configuring the AJP Connector in server.xml	262
Setting the workers.properties File	264
Configuration Settings for Apache	267
<b>Tomcat Load Balancing with Apache</b>	<b>273</b>
<b>Summary</b>	<b>282</b>
<b>Chapter 14: Tomcat and IIS</b>	<b>285</b>
<b>Concepts</b>	<b>286</b>
<b>Configuring IIS for Tomcat Out-of-Process</b>	<b>287</b>
<b>Adding Your Own Web Applications</b>	<b>296</b>
<b>Scalable Architectures with IIS and Tomcat</b>	<b>297</b>
<b>Running Tomcat In-Process</b>	<b>299</b>
<b>Log Files</b>	<b>302</b>
<b>Performance Tuning</b>	<b>304</b>
<b>Summary</b>	<b>308</b>



<b>Chapter 15: JDBC Connectivity</b>	<b>311</b>
<b>JDBC Basics</b>	<b>312</b>
Basic JDBC Operations	313
<b>Tomcat and JDBC</b>	<b>317</b>
Web Containers and RDBMSs	318
<b>Preferred Configuration: JNDI Resources</b>	<b>319</b>
<b>Alternative JDBC Configuration</b>	<b>328</b>
<b>Summary</b>	<b>332</b>
<b>Chapter 16: Tomcat Security</b>	<b>335</b>
<b>Some Basic Security Considerations</b>	<b>335</b>
<b>Securing the Filesystem</b>	<b>337</b>
Users, Groups, and Permissions	337
Recommended File Security Practices	343
<b>Securing Tomcat's Permissions</b>	<b>344</b>
<b>The Java Security Manager</b>	<b>344</b>
Overview of the Security Manager	344
Using the Security Manager with Tomcat	348
Recommended Security Manager Practices	351
<b>Security Realms</b>	<b>354</b>
Message Digests	354
Users and Roles	355
File-Based Realms	357
JDBC Realms	360
<b>SSL</b>	<b>365</b>
JSSE	366
PureTLS	370
Protecting Resources with SSL	372
<b>SSL with Apache</b>	<b>373</b>
<b>Summary</b>	<b>378</b>
<b>Chapter 17: Additional Uses for Ant</b>	<b>381</b>
<b>Introduction To Ant</b>	<b>382</b>
<b>Ant Build Process</b>	<b>384</b>
<b>Ant Build Status – E-Mail Notifications</b>	<b>389</b>

---

<b>Tomcat Ant Tasks</b>	<b>392</b>
Install	394
List	396
Stop	397
Start	398
Reload	399
Remove	399
Deploy	400
Undeploy	401
Resources	402
Roles	403
<b>Summary</b>	<b>403</b>
<hr/> <b>Chapter 18: Log4J</b>	<hr/> <b>405</b>
<b>Log4J</b>	<b>405</b>
Loggers	406
Levels	406
Appenders	407
Layouts	407
Configurators	407
Hierarchies	410
<b>Log4J in a Web Application</b>	<b>411</b>
Configuring a Logger in a Web Application	411
Logging To Files	413
Logging To the Console	416
Logging To Multiple Destinations	417
<b>Summary</b>	<b>422</b>
<hr/> <b>Chapter 19: Shared Tomcat Hosting</b>	<hr/> <b>425</b>
<b>Virtual Hosting</b>	<b>425</b>
IP-Based Virtual Hosting	426
Name-Based Virtual Hosting	428
<b>Virtual Hosting with Tomcat</b>	<b>430</b>
<b>Virtual Hosting with Tomcat 3.3</b>	<b>432</b>
Tomcat 3.3 As a Standalone Server	433
Tomcat 3.3 with Apache	436
<b>Virtual Hosting with Tomcat 4.x</b>	<b>441</b>
Tomcat 4.x As a Standalone Server	441
Tomcat 4.0 with Apache	444
<b>Fine-Tuning Shared Hosting</b>	<b>453</b>
Separate JVM for Each Virtual Host	453
Setting Memory Limits To the Tomcat JVM	457
<b>Summary</b>	<b>459</b>

# Table of Contents

---

<b>Chapter 20 Server Load Testing</b>	<b>461</b>
<b>Elements of Scalability</b>	<b>461</b>
Software Configuration	462
Deployment Architecture	464
Application Code	465
<b>Load Testing with JMeter</b>	<b>465</b>
Installing and Running JMeter	465
Making and Understanding Test Plans	466
JMeter Features	470
Interpreting Test Results	482
Distributed Load Testing	484
<b>Server Load Testing vs. Application Load Testing</b>	<b>485</b>
<b>Summary</b>	<b>487</b>
<b>Appendix A: Axis</b>	<b>491</b>
<b>Installing Axis</b>	<b>491</b>
<b>Implementing Axis</b>	<b>493</b>
<b>Appendix B: Apache SSL Setup</b>	<b>497</b>
<b>SSL and Apache</b>	<b>498</b>
Building and Installing the OpenSSL Library	499
Building and Installing mod_ssl for Apache 2.0	502
Building and Installing mod_ssl for Apache 1.3	503
Basic SSL Configuration	505
Installing a Private Key	507
Creating a Certificate Request and Temporary Certificate	508
Getting a Signed Certificate	510
<b>Summary</b>	<b>511</b>
<b>Index</b>	<b>513</b>





# 1

## Apache and Jakarta Tomcat

If you've written any Java servlets or JavaServer Pages (JSPs), chances are that you've downloaded Tomcat. That's because Tomcat is a free, feature-complete servlet container that servlet and JSP developers can use to test their code. Tomcat is also Sun's reference implementation of a servlet container, which means that Tomcat's first goal is to be 100% compliant with the versions of the Servlet and JSP specification that it supports.

However, Tomcat is more than just a test server: many individuals and corporations are using Tomcat in production environments because it has proven to be quite stable. Indeed, Tomcat is considered by many to be a worthy addition to the excellent Apache suite of products.

Despite Tomcat's popularity, it suffers from a common shortcoming among open source projects: lack of complete documentation. There is some documentation distributed with Tomcat (mirrored at <http://jakarta.apache.org/tomcat/tomcat-4.0-doc/>) and there's even an open source effort to write a Tomcat book (<http://tomcatbook.sourceforge.net/>). Even with these resources, however, there is much room for additional material.

We've created this book to fill in some of the documentation holes and use the combined experience of the authors to help Java developers and system administrators make the most of the Tomcat product. Whether you're looking to learn enough to just get started developing servlets or trying to understand the more arcane aspects of Tomcat configuration, you should find what you're looking for within these pages.

The first two chapters are designed to provide newcomers with some basic background information that will become prerequisite learning for future chapters. If you're a system administrator with no previous Java experience, you are advised to read them; likewise if you're a Java developer who is new to Tomcat. Finally, if you're well informed about Tomcat and Java, you'll probably want to jump straight ahead to Chapter 3, although skimming this chapter and its successor is likely to yield some additions to your present understanding.

We will cover the following points in this chapter:

- ❑ The origins of the Tomcat server
- ❑ The terms of Tomcat's license and how it compares to other open source licenses
- ❑ How Tomcat fits into the Java big picture
- ❑ How Tomcat can be integrated with Apache and other web servers

## Humble Beginnings: The Apache Project

One of the earliest web servers was developed by Rob McCool at the National Center for Supercomputer Applications, University of Illinois, Urbana-Champaign, referred to colloquially as the NCSA project, or NCSA for short. In 1995, the NCSA server was quite popular, but its future was uncertain as Rob left NCSA in 1994. A group of developers got together and compiled all the NCSA bug fixes and enhancements they had found and patched them into the NCSA code base. The developers released this new version in April 1995, and called it Apache, which was a sort of acronym for "A PAtCHy Web Server".

Apache was readily accepted by the web-serving community from its earliest days, and less than a year after its release it unseated NCSA to become the most used web server in the world (measured by the total number of servers running Apache), a distinction that it has held ever since (according to Apache's web site). Incidentally, during the same period that Apache's use spread, NCSA's popularity plummeted and by 1999 was officially discontinued by its maintainers.

*For more information on the history of Apache and its developers, see [http://httpd.apache.org/ABOUT\\_APACHE.html](http://httpd.apache.org/ABOUT_APACHE.html).*

Today the Apache web server is available on pretty much any major operating system – as of this writing, downloads are available for 29 different operating systems. Apache can be found running on the some of the largest server farms in the world as well as on some of the smallest devices (including the Linux-based Sharp Zaurus hand-held). In Unix data centers, Apache is as ubiquitous as air conditioning and UPS systems.

While Apache was originally a somewhat mangy collection of miscellaneous patches, today's versions are state-of-the-art, incorporating rock-solid stability with bleeding edge features. The only real competitor to Apache in terms of market share and feature set is Microsoft's Internet Information Services (IIS), which is bundled free with certain versions of the Windows operating system. At the time of writing, Apache's market share was estimated at around 56%, with IIS at a distant 32% (statistics courtesy of <http://www.netcraft.com/survey/>, June 2002).

It is also worth noting that Apache has a reputation of being much more secure than Microsoft IIS. When new vulnerabilities are discovered in either server, the Apache developers fix Apache far faster than Microsoft fixes IIS.

## The Apache Software Foundation

In 1999, the same folks who wrote the Apache server formed the Apache Software Foundation (ASF). The ASF is a non-profit organization created to facilitate the development of open source software projects. According to their web site, the ASF accomplishes this goal by:

- ❑ Providing a **foundation** for open, collaborative software development projects by supplying hardware, communication, and business infrastructure
- ❑ Creating an independent legal entity to which companies and individuals can **donate resources** and be assured that those resources will be used for the public benefit
- ❑ Providing a means for individual volunteers to be sheltered from **legal suits** directed at the Foundation's projects
- ❑ Protecting the Apache **brand**, as applied to its software products, from being abused by other organizations

In practice, the ASF does indeed sponsor a great many open source projects. While the best known of these projects is likely the aforementioned Apache web server, the ASF hosts many other well-respected and widely used projects.

## Apache Projects

The following is a listing of the current Apache projects, all of which can be found at <http://www.apache.org/>:

Project Name	Description
HTTP (Web) Server	The famous Apache web server.
Apache Portable Runtime (APR)	A library of platform-independent C code that forms a portability layer for compiling applications on multiple platforms, such as Linux, Windows, BeOS, and OS/2.
Jakarta	Apache's Java-related efforts; described in detail later in this chapter.
Perl	Umbrella for Apache's well-known <code>mod_perl</code> project, as well as other Apache-specific Perl modules. <code>mod_perl</code> enables highly efficient integration with Apache and Perl programs.
PHP	PHP is a <i>very</i> popular scripting language for dynamically creating HTML or performing business logic in HTML pages. Like <code>mod_perl</code> , PHP can be efficiently integrated with Apache.
TCL	Umbrella for Apache's efforts to tightly integrate TCL with the Apache web server (like <code>mod_perl</code> ).
XML	Umbrella for Apache's cross-platform XML projects, such as the Xerces and Crimson XML parsers and the AXIS SOAP engine.



# The Jakarta Project

Of most relevance to this book is Apache's **Jakarta** project, of which the Tomcat server is a subproject. The Jakarta project is the umbrella under which the ASF sponsors the development of Java subprojects. At the time of writing, there is an impressive array of more than twenty of these. They are divided into three different categories: "Libraries, Tools, and APIs", "Frameworks and Engines", and "Server Applications". We will highlight two projects from the first category (Ant and Log4J), one from the framework category (Struts), and, of course, Tomcat.

## Tomcat

The Jakarta Tomcat project has its origins in the earliest days of Java's servlet technology. Servlets plug into special web servers, called servlet containers (originally called servlet engines). Sun created the first servlet container, called the Java Web Server, which demonstrated the technology but wasn't terribly robust. Meanwhile, the ASF folks created the JServ product, which was a servlet engine that integrated with the Apache web server.

In 1999, Sun donated their servlet container code to the ASF, and the two projects were merged to create the Tomcat server. Today, Tomcat serves as Sun's official reference implementation (RI), which means that Tomcat's first priority is to be fully compliant with the Servlet and JSP specifications published by Sun. JSP pages are simply an alternative, HTML-like way to write servlets. We will discuss all this in more detail in the next chapter.

A reference implementation also has the side benefit of honing the specification. As developers seek to put in code that has been defined in the specifications, problems in implementation requirements and conflicts within the specifications are highlighted.

A reference implementation is in principal completely specification-compliant and therefore can be very valuable, especially for people who are using very advanced parts of the specification. The reference implementation is available at the same time as the public release of the specifications, which means that Tomcat is usually the first server out there that provides the enhanced specification features when a new specification version is completed.

The first version of Tomcat was the 3.x series, and it served as the reference implementation of the Servlet 2.2 and JSP 1.1 specifications. The Tomcat 3.x series was descended from the original code that Sun provided to the ASF in 1999.

In 2001, Tomcat 4.0 (codenamed Catalina) was released, and was a complete redesign of the Tomcat architecture and had a new code base. The Tomcat 4.x series, which is current as of this writing, is the reference implementation of the Servlet 2.3 and JSP 1.2 specifications.

At the time of writing, the latest stable version is 4.0.4. Hints of Tomcat 5.0 are on the horizon, as the new Servlet 2.4 and JSP 2.0 specifications are nearing release and Tomcat 5.0 will need to implement those specifications.

## Ant

Ant is a tool to automate building and deploying applications that range from the very simple to the extremely complex. If you're familiar with Unix, you might think this sounds like the ubiquitous `make` tool. In fact, Ant was created by a group of people who wanted to create a replacement for `make`. You can read about their comments on the subject at <http://jakarta.apache.org/ant/>.

Ant can be used for building applications in any language, and it can be used on any platform that has a Java 1.1 virtual machine or better. Ant's versatility can also be extended with Java plug-ins. Ant won awards from both the Software Development and Java World magazines in 2002, and it is extremely popular amongst developers.

## **Log4J**

Developers generally use logging for two purposes: debugging during development and monitoring when the system is in production. When developing systems, developers usually prefer logging to be as verbose as possible, and aren't concerned with its impact on the system's overall performance. However, when a system is deployed into production, developers want logging to impact performance as little as possible.

Log4J represents more than five years of work towards creating the ideal logging solution for Java programs, combining the desire for generation of rich data at development time with the need for minimal performance degradation in production environments. If your current logging technique is executing something like `System.out.println()`, you owe it to yourself to investigate this project and see what else is possible with logging.

### **Log4J Versus JDK 1.4 Logging**

Java 1.4 introduced a logging mechanism to Java as part of the standard J2SE platform. Log4J has been in its present form since late 1999, and thus predates the JDK 1.4 logging mechanism by a little more than 2 years (JDK 1.4 went final in early 2002). When it was learned that Java 1.4 would incorporate logging, the Log4J group lobbied to have its product incorporated into Java as the official logging mechanism for the platform. However, that did not happen.

With the release of Java 1.4, Log4J didn't disappear, and doesn't intend to. Log4J provides two advantages over the Java 1.4 logging mechanism: it has more features and it can be used with Java 1.1 or later.

## **Struts**

The current architectural best practice for web applications is the Model View Controller (MVC) design pattern. Under this model, the application is divided into three logical layers (also called tiers): the View, which represents the user interface; the Model, which represents the business logic specific to the application including any persistent data store (for example, a database); and the Controller, which coordinates how the View and the Model interact, and takes care of any other general application behavior (for example, application lifecycle issues). We'll see more on the MVC architecture in the next chapter.

Servlets and JavaServer Pages are the standard Java way to create web applications. They provide an efficient interface to the Web's HTTP protocol. However, developers who wish to create an MVC architecture with servlets and JSP must still do quite a bit of work.

Many third-party frameworks have been created which attempt to relieve developers from the burden of implementing their own MVC architecture, freeing them to instead focus on solving the unique business problems of their organization. Struts is one of these frameworks. Struts has gained an excellent reputation in the development community as being well-designed and very flexible.

### **Other Jakarta Subprojects**

There are many other Jakarta subprojects, including: Lucene, a full-featured search engine; Jetspeed, a portal server; and James, a mail server. See these and others at <http://jakarta.apache.org/>.

# Distributing Tomcat

Tomcat is open source software, and as such is free and freely distributable. However, if you have much experience in dealing with open source software, you're probably aware that the terms of distribution can vary from project to project.

Most open source software is released with an accompanying license that states what may and may not be done to the software. There are at least forty different open source licenses out there, each of which has slightly different terms.

Providing a primer on all of the various open source licenses is beyond the scope of this chapter, but the license governing Tomcat will be discussed here and compared with a few of the more popular open source licenses.

Tomcat is distributed under the Apache License, which can be read from the `$CATALINA_HOME/LICENSE` file. The key points of this license state that:

- ❑ The Apache License must be included with any redistributions of Tomcat's sourcecode or binaries
- ❑ Any documentation included with a redistribution must give a nod to the ASF
- ❑ Products derived from the Tomcat sourcecode can't use the terms "Tomcat", "The Jakarta Project", "Apache", or "Apache Software Foundation" to endorse or promote their software without prior written permission from the ASF
- ❑ Tomcat has no warranty of any kind

However, through omission, the license contains these additional implicit permissions:

- ❑ Tomcat can be used by any entity, commercial or non-commercial, for free without limitation
- ❑ Those who make modifications to Tomcat and distribute their modified version do not have to include the sourcecode of their modifications
- ❑ Those who make modifications to Tomcat do not have to donate their modifications back to the ASF

Thus, you're free to deploy Tomcat in your company in any way you see fit. It can be your production web server or your test servlet container used by your developers. You can also redistribute Tomcat with any commercial application that you may be selling, provided that you include the license and give credit to the ASF. You can even use the Tomcat sourcecode as the foundation for your own commercial product

## Comparison with Other Licenses

Among the previously mentioned rather large group of other open source licenses, there are two licenses which are particularly popular at the present time: the GNU General Public License (GPL) and the GNU Lesser General Public License (LGPL). Let's take a look at how each of these licenses compare to the Apache License.

## GPL

The GNU Project created and actively evangelizes the GPL. The GNU Project is somewhat similar to the ASF, with the exception that the GNU Project would like all of the non-free (that is, closed source or proprietary) software in the world to become free; the ASF has no (stated) desire to do this and simply wants to provide free software.

### **What Does It Mean to Be Free?**

Free software can mean one of two entirely different things: software that doesn't cost anything, and software that can be freely copied, distributed, and modified by anyone (thus the sourcecode is included or available); such software can be distributed either for free or for a fee. A simpler way to explain the difference between these two types of free is "free as in free beer" and "free as in free speech". The GNU Project's goal is to create free software of the latter category. All uses of the phrase "free software" in the remainder of this section will use this definition.

The differences between the Apache License and the GPL thus mirror the distinct philosophies of the two organizations. Specifically, the GPL has these key differences from the Apache License:

- ❑ No non-free software may contain GPL-licensed products or use GPL-licensed sourcecode. If non-free software is found to contain GPL-licensed binaries or code, it must remove such elements or become free software itself.
- ❑ All modifications made to GPL-licensed products must be released as free software if the modifications are also publicly released.

These two differences have huge implications for commercial enterprises. If Tomcat were licensed under the GPL, any product that contained Tomcat would also have to be free software.

Furthermore, while the Apache License permits an organization to make modifications to Tomcat and sell it under a different name as a closed source product, the GPL would not allow any such act to occur; the new derived product would also have to be released as free software.

## LGPL

The LGPL is similar to the GPL, with one major difference: non-free software may contain LGPL-licensed products. The LGPL license is intended primarily for software libraries that are themselves free software but whose authors want them to be available for use by companies who produce non-free software.

If Tomcat were licensed under the LGPL, it could be embedded in non-free software, but Tomcat could not itself be modified and released as a non-free software product.

*For more information on the GPL and LGPL licenses, see <http://www.gnu.org/>.*

## Other Licenses

Understanding and comparing open source licenses can be a rather complex task. The explanations above are an attempt to simplify the issues. For more detailed information on these and other licenses, there are two specific resources that can help you:

- ❑ The Open Source Initiative (OSI) maintains a database of open source licenses. Visit them at <http://www.opensource.org/>.
- ❑ The GNU Project, mentioned above, has an extensive comparison of open source licenses with the GPL license. See it at <http://www.gnu.org/licenses/license-list.html>.

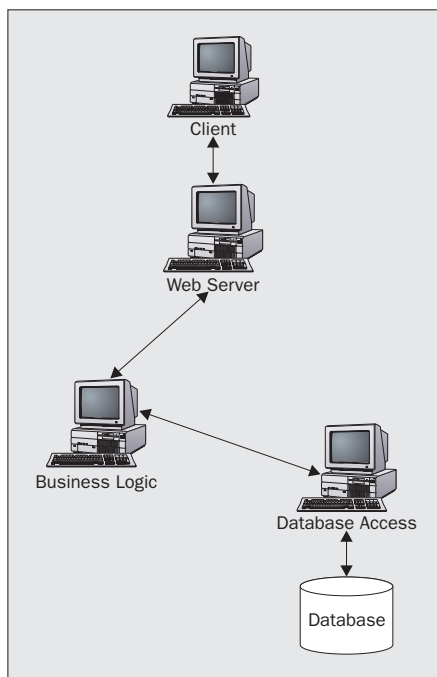
## The Big Picture: J2EE

As a servlet container, Tomcat is a key component of a larger set of standards collectively referred to as the Java 2 Platform, Enterprise Edition (J2EE). J2EE defines a group of Java-based code libraries (called APIs in the Java world) that are suited to creating web applications for the enterprise (that is, a large company). To be sure, companies of any size can take advantage of the J2EE technologies, but J2EE is especially designed to solve the problems associated with the creation of large software systems. Java developers can download all of the J2EE APIs in a single ZIP archive, which comes complete with binaries and documentation.

## Distributed Systems

We saw in the previous section that J2EE is designed with the creation of large software systems in mind. You, the reader, may ask, "What is so different about large software systems versus small systems?" The answer lies in the notion of **distributed systems**.

A distributed system is one in which the various components of the system are distributed across multiple different machines. For example, in a given web application, one server might handle receiving and responding to HTTP requests while another server handles all the business logic, and another server handles all the database I/O:



So why create a distributed system? Large web applications and enterprise systems can have enormous demands placed upon them. Frequently, these demands are larger than a single server can meet. While one may be tempted to simply upgrade the server, adding more processors and more memory to it, these upgrades can only stretch the server's capacity so far. Every server, no matter how expandable and efficient it may be, will run up against limitations.

In these situations, where a single server simply cannot meet the needs of a software system, a distributed system is ideal. Multiple servers can work together to meet demands that a single server could not.

But, why bother creating a distributed system? Why not just copy the same application and deploy it on multiple servers? Surely load balancing – splitting up the requests amongst the different servers – can then handle the load?

It turns out that by separating the different logical components of an application, system administrators are free to tune the various components of the distributed system according to their unique needs. For example, the servers that read and write to the database will require different optimizations than those servers which are communicating with web clients via HTTP over TCP.

## The J2EE APIs

Creating the distributed systems that we've described in the preceding paragraphs is not an easy task. We mentioned a few paragraphs ago that J2EE is a collection of code libraries called APIs. The J2EE APIs are designed to work together to simplify the creation of robust and efficient distributed systems. Here is a list of some of the key J2EE technologies and a brief description of each:

J2EE API	Description
Enterprise JavaBeans (EJB)	EJB technology provides a simple mechanism for creating distributed business logic components. EJB authors follow a simple pattern to write business logic and the rest of the low-level details relating to lifecycle, distribution, persistence, and so on are handled automatically.
Java Message Service (JMS)	JMS provides asynchronous messaging capability to J2EE applications.
Java Naming and Directory Service (JNDI)	JNDI enables J2EE applications to communicate with registries and directories. A registry/directory is a centralized location for storing business information. JNDI supports the industry standard LDAP protocol and interfaces with many other popular registry standards. JNDI makes it possible to centralize the configuration of a distributed system.
Servlets	As explained earlier in this chapter, servlets work with special servers called servlet containers to process HTTP requests and send HTTP responses. Servlets often work directly with EJBs.
JavaServer Pages (JSP)	JSP technology is an alternative, HTML-like interface for creating servlets. At runtime, the servlet container converts a JSP into a servlet.

*Table continued on following page*