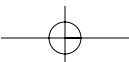
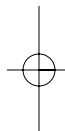
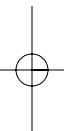


# **Effective Enterprise Java**



# *Effective* **SOFTWARE DEVELOPMENT SERIES**

Scott Meyers, Consulting Editor

The **Effective Software Development Series** provides expert advice on all aspects of modern software development. Books in the series are well written, technically sound, of lasting value, and tractable length. Each describes the critical things the experts almost always do—or almost always avoid doing—to produce outstanding software.

Scott Meyers (author of the *Effective C++* books and CD) conceived of the series and acts as its consulting editor. Authors in the series work with Meyers and with Addison-Wesley Professional's editorial staff to create essential reading for software developers of every stripe.

## **TITLES IN THE SERIES**

Elliott Rusty Harold, *Effective XML: 50 Specific Ways to Improve Your XML*  
0321150406

Diomidis Spinellis, *Code Reading: The Open Source Perspective* 0201799405

---

For more information on books in this series please see [www.awprofessional.com/esds](http://www.awprofessional.com/esds)

# Effective Enterprise Java

Ted Neward

◆ Addison-Wesley

Boston • San Francisco • New York • Toronto • Montreal  
London • Munich • Paris • Madrid  
Capetown • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Addison-Wesley was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers discounts on this book when ordered in quantity for bulk purchases and special sales. For more information, please contact:

U.S. Corporate and Government Sales  
(800) 382-3419  
corpsales@pearsontechgroup.com

For sales outside of the U.S., please contact:

International Sales  
(317) 581-3793  
international@pearsontechgroup.com

Visit Addison-Wesley on the Web: [www.awprofessional.com](http://www.awprofessional.com)

***Library of Congress Cataloging-in-Publication Data***

Neward, Ted.

Effective Enterprise Java / Ted Neward.

p. cm.

ISBN 0-321-13000-6 (pbk. : alk. paper)

1. Java (Computer program language) I. Title.

QA76.73.J38N48 2004

005.13'3—dc22

2004012164

Copyright © 2005 by Pearson Education, Inc.

All rights reserved. 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, or otherwise, without the prior consent of the publisher. Printed in the United States of America. Published simultaneously in Canada.

For information on obtaining permission for use of material from this work, please submit a written request to:

Pearson Education, Inc.  
Rights and Contracts Department  
75 Arlington Street, Suite 300  
Boston, MA 02116  
Fax: (617) 848-7047

ISBN 0-321-13000-6

Text printed on recycled paper

1 2 3 4 5 6 7 8 9 10—CRS—0807060504

First printing, August 2004

# Contents

	<b>Foreword</b>	<b>ix</b>
	<b>Preface</b>	<b>xi</b>
	About the items	xiii
	Acknowledgments	xiv
	Reporting bugs, making suggestions, and getting book updates	xvi
	<b>List of Abbreviations</b>	<b>xvii</b>
<b>Chapter 1</b>	<b>Introduction</b>	<b>1</b>
	The goals of J2EE	3
	Middleware and J2EE	5
	J2EE implementation	10
	The ten fallacies of enterprise computing	14
<b>Chapter 2</b>	<b>Architecture</b>	<b>19</b>
	Item 1: Prefer components as the key element of development, deployment, and reuse	19
	Item 2: Prefer loose coupling across component boundaries	26
	Item 3: Differentiate layers from tiers	31
	Item 4: Keep data and processors close together	35
	Item 5: Remember that identity breeds contention	40
	Item 6: Use hook points to inject optimizations, customizations, or new functionality	47
	Item 7: Be robust in the face of failure	53
	Item 8: Define your performance and scalability goals	59
	Item 9: Restrict EJB to transactional processing	64
	Item 10: Never optimize without profiling first	67
	Item 11: Recognize the cost of vendor neutrality	74
	Item 12: Build in monitoring support	78
	Item 13: Build in administration support	87
	Item 14: Make deployment as simple as possible	96

<b>Chapter 3</b>	<b>Communication</b>	<b>101</b>
	Item 15: Understand all your communications options	101
	Item 16: Consider your lookup carefully	108
	Item 17: Recognize the cost of network access	114
	Item 18: Prefer context-complete communication styles	120
	Item 19: Prefer data-driven communication over behavior-driven communication	128
	Item 20: Avoid waiting for remote service requests to respond	136
	Item 21: Consider partitioning components to avoid excessive load on any one machine	140
	Item 22: Consider using Web Services for open integration	147
	Item 23: Pass data in bulk	150
	Item 24: Consider rolling your own communication proxies	155
<b>Chapter 4</b>	<b>Processing</b>	<b>159</b>
	Item 25: Keep it simple	160
	Item 26: Prefer rules engines for complex state evaluation and execution	163
	Item 27: Prefer transactional processing for implicitly nonatomic failure scenarios	169
	Item 28: Differentiate user transactions from system transactions	175
	Item 29: Minimize lock windows	180
	Item 30: Never cede control outside your component while holding locks	187
	Item 31: Understand EJB transactional affinity	193
	Item 32: Prefer local transactions to distributed ones	196
	Item 33: Consider using optimistic concurrency for better scalability	200
	Item 34: Consider using pessimistic concurrency for explicit concurrency control	206
	Item 35: Consider lower isolation levels for better transactional throughput	211
	Item 36: Use savepoints to keep partial work in the face of rollback	216
	Item 37: Replicate resources when possible to avoid lock regions	219
	Item 38: Favor the immutable, for it needs no locks	222
<b>Chapter 5</b>	<b>State Management</b>	<b>225</b>
	Item 39: Use <code>HttpSession</code> sparingly	227
	Item 40: Use objects-first persistence to preserve your domain model	232
	Item 41: Use relational-first persistence to expose the power of the relational model	236
	Item 42: Use procedural-first persistence to create an encapsulation layer	246
	Item 43: Recognize the object-hierarchical impedance mismatch	249

Item 44:	Use in-process or local storage to avoid the network	259
Item 45:	Never assume you own the data or the database	263
Item 46:	Lazy-load infrequently used data	266
Item 47:	Eager-load frequently used data	271
Item 48:	Batch SQL work to avoid round-trips	274
Item 49:	Know your JDBC provider	277
Item 50:	Tune your SQL	281
<b>Chapter 6</b>	<b>Presentation</b>	<b>285</b>
Item 51:	Consider rich-client UI technologies	286
	Dynamic HTML	290
	Macromedia Flash	290
	Applets	291
	The <code>URLClassLoader</code> class	292
	Java Network Launch Protocol (JNLP) and Java Web Start	293
Item 52:	Keep HTML minimal	296
Item 53:	Separate presentation from processing	299
Item 54:	Keep style separate from content	305
Item 55:	Pregenerate content to minimize processing	309
Item 56:	Validate early, validate everything	311
<b>Chapter 7</b>	<b>Security</b>	<b>321</b>
Item 57:	Security is a process, not a product	325
Item 58:	Remember that security is not just prevention	329
Item 59:	Establish a threat model	333
Item 60:	Assume insecurity	336
Item 61:	Always validate user input	342
Item 62:	Turn on platform security	352
Item 63:	Use role-based authorization	356
Item 64:	Use <code>SignedObject</code> to provide integrity of Serialized objects	365
Item 65:	Use <code>SealedObject</code> to provide confidentiality of Serializable objects	370
Item 66:	Use <code>GuardedObject</code> to provide access control on objects	372
<b>Chapter 8</b>	<b>System</b>	<b>379</b>
Item 67:	Aggressively release resources	379
Item 68:	Tune the JVM	387
Item 69:	Use independent JREs for side-by-side versioning	395
Item 70:	Recognize <code>ClassLoader</code> boundaries	400
	Isolation	403
	Versioning	404

**viii** | *Contents*

Item 71: Understand Java Object Serialization	408
The <code>serialVersionUID</code> field	411
Customization ( <code>writeObject</code> and <code>readObject</code> )	412
Replacement ( <code>writeReplace</code> and <code>readResolve</code> )	413
Further Details	415
Item 72: Don't fight the garbage collector	416
Item 73: Prefer container-managed resource management	426
Item 74: Use reference objects to augment garbage collection behavior	430
<code>SoftReference</code> objects	434
<code>WeakReference</code> objects	437
<code>PhantomReference</code> objects	439
Item 75: Don't be afraid of JNI code on the server	445
<b>Bibliography</b>	<b>449</b>
<b>Index</b>	<b>457</b>



# Foreword

Designing and implementing large-scale enterprise systems is hard. Building effective enterprise Java deployments is even harder. I see these difficulties on a daily basis. When consulting on enterprise projects, I see the real-world issues that developers are facing. I have also seen discussions, frustrations, and solutions to some of the issues on a daily basis on TheServerSide.com (Your Enterprise Java Community). TheServerSide.com really grew from the needs of developers faced with the new world of J2EE. It was the water cooler that allowed us to chat about solutions that worked for us, and it saw the growth of enterprise Java patterns.

Developing for the enterprise is a very different beast when compared to building smaller, standalone applications. We have to consider issues that we can safely ignore in the other world. As soon as we have to *share data* among multiple users, we start down the enterprise path. Then we start facing questions: What is the best solution for allowing concurrency to this data? How coherent and correct does it have to be? How can we scale up from 2 to 50 to 1,000 clients? These are all significant questions, and I don't feel that the average developer has enough help in answering them. Well, simply answering the questions may not be the correct focus. We need to be taught about the various issues involved and shown techniques that can help with the various problems. With Ted Neward's book, we are now armed with the knowledge that will allow us to come up with the right balance in the solution for each particular problem.

No book has attacked these problems quite like *Effective Enterprise Java* does. The most important part of this book is that it teaches you two things really well.

**You will understand the general issues of enterprise computing.**

These enterprise problems are far from new. Ted has been around the block, and he understands the core issues at work. A non-Java developer would get a lot out of this book for this very reason. What you learn here

will be with you for as long as you develop enterprise solutions. The language and APIs may change, but you will understand the issues in building a good architecture, the options you have for communication, the choices for where to store state, the various security concerns, and so much more.

**You will be able to attack the problems by using enterprise Java.**

Although the book offers genuine insight into the general enterprise problems, it also gives you tools to solve them with enterprise Java today. You will understand more about where the various enterprise Java technologies fit together. When would you use Web Services? What can messaging do for you? What is EJB good for? This book provides answers to these questions.

It is great to have some answers to these common questions. The style of the book, in which you are given a set of “effective items,” gets right to the point. Get stuck in, and enjoy the ride!

Dion Almaer  
Editor-in-Chief, [TheServerSide.com](http://TheServerSide.com)

# Preface

*Those who cannot remember the past are doomed to repeat it.*

—George Santayana

These are heady days for Java programmers. Commercially available for less than a decade, Java has nevertheless emerged as the language of choice for enterprise systems on nearly all major computing platforms. Companies and individuals with challenging programming problems increasingly embrace the language and platform, and the question faced by those who do not use Java is often *when* they will start, not *if*. The breadth and scope of the specifications and libraries defined by and under the Java 2 Enterprise Edition Specification—which both dwarfs and subsumes that of the Java language itself—makes it possible to write rich, complex systems without sacrificing performance or implementing common algorithms or data structures from scratch. The Java language and virtual machines grow ever more powerful. Tools and environments for Java development grow ever more abundant and robust. Commercial libraries all but obviate the need to write code in many application areas.

Scott Meyers wrote much the same in his opening to *More Effective C++* [Meyers97] almost a decade ago. It seems a fitting tribute to use that paragraph, suitably modified, as the opening in this book. As a matter of fact, the two paragraphs are deliberately side-by-side similar. In many ways, we now find ourselves in a part of the Golden Age of Java, looking out over a landscape that stretches from horizon to horizon, with ample space and established borders that seem to have no limit. Just as C++ ruled the landscape in 1996, Java rules the landscape in 2004.

The chief aim in drawing these parallels is to recognize the scenario—not more than two years after Scott wrote that paragraph, C++ was toppled from its throne by this upstart named Java. Just as C++ developers had

finally begun to “figure everything out,” this new language and environment leapt forward into the fray, and almost overnight, it seemed, took over. In turn, Java now faces fierce competition from Microsoft’s .NET platform. The natural concern, then, is to see that history doesn’t repeat itself. To do that, Java developers must make sure that the systems they develop meet or exceed expectations set for them. To do *that*, Java developers need to know how to make the most of the language and platform they use.

It has been said, in many places by many people, that it takes about five years to “figure out” a technology and how best to use it. Certainly this was true for C++: in 1990, we looked at C++ as simply another object-oriented language, and therefore using it should mirror the best-usage practices discovered with Smalltalk. By 1995, we were well out of that world and starting to explore the uniqueness that C++ provides (such as templates and the STL). Certainly this was also true of HTTP: in 1995, when the browser debuted, we looked at HTTP as the means by which HTML was delivered. Now, we look at HTTP as a universal transport by which to transmit all sorts of data.

Thus, the timing is fortuitous for Java. It officially debuted in 1995; for all intents and purposes, however, Java truly entered the mindscape of the average developer in 1997 or so, having by that point built enough of a “critical mass” to win its way past the critics and skeptics. Almost a decade later, we have been writing Java applications for most of that time, and we’re starting to see the practices and patterns that have emerged to assist (but not necessarily guarantee) successful deployments. As a community, we’re just starting to hit our stride.

Some things aren’t different from the C++ days. We have the same questions, modified for the Java world, as those we asked (and Scott answered) a decade ago about C++. As the language and platform have matured and our experience with Java has increased, our needs for information have changed. In 1996, people wanted to know *what* Java was. “It has something to do with the Internet, whatever that is” was a common explanation. Initially, developers focused on using applets, making rich browser clients, and harnessing Java’s cross-platform portability. By 1998, they wanted to know *how* to make it work: “How do I access a relational database? How do I internationalize? How do I reach across physical machine boundaries?” Now, Java programmers ask higher-level questions: “How can I design my enterprise systems so they will adapt to future demands?”

How can I improve the efficiency of my code without compromising its correctness or making it harder to use? How can I implement sophisticated functionality not directly supported by the language or platform?”

As if this weren't enough, a new dimension has arisen in the whole enterprise system arena, neatly captured in two words: Web Services. Even as Java developers are asking the hard higher-level questions about Java, they face the start of the cycle with Web Services—what is a Web Service, how does it work, and, perhaps most importantly, how does it relate to Java?

In this book, I answer these questions and many like them.

---

### About the items

One thing I feel compelled to point out, before we get too deeply into it all, is that readers may notice a significant difference between the items in this book and those from books like *Effective Java* [Bloch] and *Effective C++* [Meyers95]. In particular, the scope of the items in this book is much larger than that in other similar books—in here, there's less focus on language and/or APIs and more on design-level constructs and/or usage patterns.

This is not an accident; in fact, I believe that this is in keeping with the larger scope of an enterprise application as a whole. Certainly, without question, all of *Effective Java* also applies to building enterprise applications, but to simply remain at that level misses the larger point, that enterprise systems have much more to worry about, things outside of the scope of the language or APIs.

For example, an unsuccessful EJB application usually begins not with misuse of a particular method call or interface but with the design of entity beans that are called directly from a client. This isn't so much an implementation problem as a design problem, and solving it requires a more “high-level” view of what the entity bean is trying to provide in general. (See Item 40 for more details about entity beans and their consequences.)

As a result, the items presented in this book strive to help developers recognize efficiency not at a language level but at a systemic and architectural level. Many of these items will be familiar ground to some; many will be simple codification of something some readers “always knew.” That's OK—what's “intuitive” to one reader will be new to another, and vice versa.

In addition, I have carefully tried to avoid walking on familiar territory. A number of books have been released that discuss best practices and effective use of the virtual machine and/or the language; see the Bibliography for a complete list. As a result, discussions of material covered there won't be repeated here unless it has some particular relevance or application within the enterprise space.

Toward that end, I will refer frequently to patterns already established within the enterprise Java literature space; in particular, I will tend to lean heavily on Fowler's *Patterns of Enterprise Application Architecture* (Addison-Wesley, 2002), Hohpe and Woolf's *Enterprise Application Integration* (Addison-Wesley, 2004), and Alur, Crupi, and Malks's *Core J2EE Patterns*, 2nd ed. (Addison-Wesley, 2003), among others. (Again, see the Bibliography for a complete list.)

Where a pattern is cited by name, I use the standard Gang-of-Four pattern citation format, citing the pattern name with its page number in parentheses behind it; however, because these patterns come from different sources, I also put the author's names (or "GOF" for the Gang-of-Four book, *Design Patterns*) as part of the page citation. So a reference to the Data Transfer Object pattern from Fowler's *Patterns of Enterprise Application Architecture* book will be cited as "Data Transfer Object ([Fowler401])".

---

## Acknowledgments

Authors have a tendency to spend a lot of time thanking people; there's a reason for that.

First and foremost, I want to thank my peers in the industry, most of all my fellow instructors at DevelopMentor. Kevin Jones, Brian Maso, Stu Halloway, Simon Horrell, Dan Weston, and Bob Beauchemin all served as sounding boards for the topics in this book at some point over its 30-month gestation. Tim Ewald, Don Box, Fritz Onion, Keith Brown, Mike Woodring, Ingo Rammer, and Peter Drayton all gave me insights into the Java platform by routinely smashing my preconceptions and assumptions in discussions that had nothing to do with Java. Outside of DevelopMentor, my fellow NoFluffJustStuff Symposium speakers Dion Alamer, Bruce Tate, Mike Clark, Erik Hatcher, Glenn Vandenburg, Dave Thomas, Jason Hunter, James Duncan Davidson, and Ron Bodkin, among others, all forced me to justify my assertions, questioned my conclusions,

and offered suggestions and tips designed to make the book a stronger work. Thanks to Jay Zimmerman for inviting me to be a part of NoFluffJustStuff in the first place. Numerous other speakers at conferences far and wide (far too many to list here) played a similar role, known in the common vernacular as “keeping me honest.”

Second, I must tip the hat to the staff at Addison-Wesley. This book took over twice as long as it was supposed to, and never once did Mike Hendrickson, the editor who started the project, nor Ann Sellers, the editor who inherited the project, use anything but polite language when asking if it was done—their patience far exceeded what mine would have been in their shoes. The reviewers, both of the content when it was hosted on the Web via my blog, as well as the more finalized manuscript drafts, did a Herculean job reading through the material and offering up copious corrections, suggestions, enhancements, and ideas. Thanks to Matt Anderson, Kevin Bentley, Dave Cooke, Mary Dageforde, Kevin Davis, Matthew P. Johnson, and Bruce Scharlau for all their help.

Writing this book has been both a labor of absolute love and an exercise in abject terror. While I’ve always looked for a project like this to let me rant about my thoughts on enterprise Java development, few books have “raised the bar” as the previous *Effective* books have done. With *Effective C++*, Scott Meyers gave me (and millions of other neophyte C++ programmers) the leg up I needed to start using C++, rather than just flailing around with it. Then Joshua Bloch wrote *Effective Java*, bringing the beauty of the item format to the Java platform. Eliote Rusty Harold further upped the ante with *Effective XML*. If ever an author were looking for an intimidating set of authors to follow, those three fill out the set admirably. Fortunately, help was available in the form of Scott Meyers, who spent almost as much time reviewing and criticizing (in the good sense) this book as I did writing it. His comments and insights helped transform a tolerable collection of suggestions into the book you see in front of you. Scott, I owe you a tremendous debt of gratitude, both for your help during the last year and your help 10 years ago as I struggled to understand C++. It has been a privilege and an honor to do this with you; thank you.

Finally, of course, I must thank my family and friends, those loving people who periodically staged interventions and dragged me, kicking and screaming, back into this thing they kept calling “the real world”: parties, holidays, even just hanging out for a night or two relaxing over the

Nintendo64 or Xbox. Though I'll never admit it out loud, they kept me sane in the face of the mounting pressure of writing this book during the last two years.

---

### Reporting bugs, making suggestions, and getting book updates

I have tried to make this book as accurate, readable, and useful as possible, but I know there is room for improvement. (There is *always* room for improvement.) If you find an error of any kind—technical, grammatical, typographical, spiritual, *whatever*—please tell me about it. I will try to ensure the mistake is corrected in a future printing of the book, and if you are the first person to report it, I will happily add your name to the book's acknowledgments. Likewise, if you have suggestions or ideas on how to improve the book for subsequent revisions or editions, I'm all eyes and ears.

I continue to collect guidelines for effective enterprise programming in Java. If you have ideas for new guidelines, I'd be delighted if you'd share them with me. You can either find me on one of several public Java programming mailing lists, the predominant one being the ADVANCED-JAVA list at DISCUSS.DEVELOP.COM, or you can reach me at:

Ted Neward  
c/o Addison-Wesley Professional/Prentice Hall PTR  
Pearson Technology Group  
75 Arlington St., Suite 300  
Boston, MA 02116

Alternatively, you can drop me an email at [ted@neward.net](mailto:ted@neward.net).

I maintain a list of changes to this book since its first printing (including bug fixes, clarifications, commentary, and technical updates) on the book's blog, <http://www.neward.net/ted/EEJ/index.jsp>. Please feel free to post comments and/or errata there if you wish to share them with your fellow readers.

Enough preliminaries. On with the show!



# List of Abbreviations

<b>ACID</b>	atomic, consistent, isolated, and durable
<b>AWT</b>	Abstract Windowing Toolkit
<b>BLOB</b>	Binary Large Object
<b>BMP</b>	Bean-Managed Persistence
<b>CMP</b>	Container-Managed Persistence
<b>COM</b>	Component Object Model
<b>CORBA</b>	Common Object Request Broker Architecture
<b>CSS</b>	Cascading Style Sheets
<b>DCOM</b>	Distributed COM
<b>DHTML</b>	Dynamic HTML
<b>DMZ</b>	demilitarized zone
<b>DNS</b>	Domain Name System
<b>DOM</b>	Document Object Model
<b>DTC\$</b>	distributed transaction controllers
<b>EJB</b>	Enterprise Java Bean
<b>HTML</b>	Hyper-Text Markup Language
<b>HTTP</b>	Hyper-Text Transmission Protocol
<b>IDL</b>	Interface Description Language
<b>IIOP</b>	Internet Inter-Orb Protocol
<b>IPC</b>	interprocess communication
<b>J2EE</b>	Java 2 Enterprise Edition
<b>J2SE</b>	Java 2 Standard Edition
<b>JAAS</b>	Java Authentication and Authorization Service
<b>JAXB</b>	Java API for XML Binding
<b>JAXM</b>	Java API for XML Messaging

**xviii** | *List of Abbreviations*

<b>JAXP</b>	Java API for XML Parsing
<b>JAX-RPC</b>	Java API for XML RPC
<b>JCA</b>	Java Connector API
<b>JDBC</b>	Java DataBase Connectivity
<b>JDK</b>	Java Development Kit
<b>JDO</b>	Java Data Objects
<b>JESS</b>	Java Expert System Shell
<b>JIT</b>	just-in-time
<b>JITA</b>	just-in-time activation
<b>JMS</b>	Java Message Service
<b>JMX</b>	Java Management Extensions
<b>JNDI</b>	Java Naming and Directory Interface
<b>JNI</b>	Java Native Interface
<b>JNLP</b>	Java Network Launch Protocol
<b>JRE</b>	Java Runtime Environment
<b>JSP</b>	Java Server Pages
<b>JSR</b>	Java Specification Request
<b>JSSE</b>	Java Secure Sockets Extension
<b>JTA</b>	Java Transaction API
<b>JVM</b>	Java Virtual Machine
<b>JVMCI</b>	Java Virtual Machine Debug Interface
<b>JVMPI</b>	Java Virtual Machine Profiler Interface
<b>JVMTI</b>	Java Virtual Machine Tools Interface
<b>LAN</b>	local area network
<b>MDBs</b>	Message-Driven Beans
<b>MIB</b>	Message Information Block
<b>MVC</b>	Model-View-Controller
<b>NAT</b>	Network Address Translation
<b>NFS</b>	Network File System
<b>OODBMS</b>	object-oriented database management system
<b>ORB</b>	Object Request Broker
<b>OSI</b>	
<b>OWASP</b>	Open Web Application Security Project
<b>POJOs</b>	plain old Java objects
<b>POP3</b>	Post Office Protocol v 3

<b>RDBMS</b>	relational database management system
<b>RMI</b>	Remote Method Invocation
<b>RMI/IIOP</b>	RMI over IIOP
<b>RMI/JRMP</b>	RMI over Java Remote Method Protocol
<b>RPC</b>	Remote Procedure Call
<b>SAX</b>	Streaming API for XML
<b>SMTP</b>	Simple Mail Transport Protocol
<b>SNMP</b>	Simple Network Management Protocol
<b>SOAP</b>	Simple Object Access Protocol
<b>SSL</b>	Secure Sockets Layer
<b>STL</b>	Standard Template Library
<b>SWT</b>	Standard Widget Toolkit
<b>TLS</b>	Transport Layer Security
<b>TPC</b>	two-phase commit
<b>TTL</b>	time-to-live value
<b>URI</b>	Universal Resource Identifier
<b>URL</b>	Universal Resource Locator
<b>URN</b>	Universal Resource Name
<b>VM</b>	virtual machine
<b>W3C</b>	World Wide Web Consortium
<b>WSDL</b>	Web Services Definition Language
<b>WS-I</b>	Web Services-Interoperability
<b>XML</b>	Extensible Markup Language
<b>XSD</b>	XML Schema Definition
<b>XSLT</b>	XSL:Transformation, commonly also written as XSL:T

