

Steal This Book!

Yes, you read that right. Steal this book. For free.

Nothing. Zero. Zilch. Nada. Zip.

Undoubtedly you're asking yourself, "Why would he give away a book he probably spent six grueling months writing? Is he crazy?"

The answer...is yes. Sort of. I know that every day you're faced with hundreds of computer titles, and a lot of them don't deliver the value or the information you need. So here's the deal: I'll give you this book (or this chapter, if you've only downloaded part of the book) for free provided you do me a couple favors:

1. **Send this book to your friends:** No, not your manager. Your "with it" computer friends who are looking for the next Big Thing. JXTA is it. Trust me. They want to know about it.
2. **Send a link to the book's web site:** Maybe the book is too big to send. After all, not everyone can have a fibre optic Internet connection installed in their bedroom. The site, at www.brendonwilson.com/projects/jxta, provides chapter-sized PDFs for easy downloading by the bandwidth-challenged.
3. **Visit the book's web site:** Being a professional developer, you probably have Carpal Tunnel Syndrome and shudder at the idea of typing in example source code. Save yourself the trouble. Go to www.brendonwilson.com/projects/jxta and download the source code. And while you're there, why not download some of the chapters you're missing?
4. **Buy the book:** You knew there had to be a catch. Sure, the book's PDFs are free, but I'm hoping that enough of you like the book so much that you have to buy a copy. Either that, or none of you can stand to read the book from a screen (or, worse yet, print it all out <shudder>) and resort to paper to save what's left of your eyesight. The book is available at your local bookstore or from Amazon.com (at a **handsome discount**, I might add).

I now return to your regularly scheduled program: enjoy the book!



Glossary

advertisement A description of a resource made available on a P2P network. Types of resources described by advertisements include peers, peer groups, pipes, endpoints, content, and services. In JXTA, advertisements are represented using XML documents.

content A generic term used to describe any type of text or binary data that can be stored and retrieved at a later time.

CVS Concurrent Versions System, an open-source version-control system that allows multiple developers to work on source files simultaneously. CVS tracks changes made to source code files and allows developers to merge changes back into a version-control system, to provide an

accurate, trackable, and reproducible record of a source file at any point in its development history.

endpoint An abstraction of a peer's underlying native network interfaces. The term is also used to describe the source peer sending a message or the destination peer receiving a message through the P2P network.

HTML Hypertext Markup Language, a set of text tags that are used to mark up plain text to provide formatting for a web browser. The HTML standard is currently migrating toward an XML-compatible schema.

HTTP Hypertext Transfer Protocol, a protocol used to transfer HTML and web content. HTTP is a

stateless protocol, which means that each request is completely independent of the requests that preceded it.

IETF Internet Engineering Task Force, a standards body responsible for defining standard protocols for the Internet. Standards are developed by the IETF with the participation of the Internet community, and they are documented in Requests for Comments (RFCs).

IM Instant messaging, a form of electronic communication that allows two or more parties to exchange text messages instantly. Examples of instant-messaging applications include ICQ, MSN Messenger, AOL Instant Messenger, and Yahoo! Messenger.

IP Internet Protocol, a protocol used to send data from a source computer to a destination computer. IP is a low-level protocol that is responsible for sending packets of data across a network, possibly using a variety of routes to a destination. Packets sent across the network are treated independently, defining IP as a connectionless protocol. Both the source and the destination for the packet are uniquely described in the packet's headers by an IP address. Unlike TCP, IP is not a reliable protocol and does not guarantee packet delivery.

JAR Java Archive, a compressed binary archive format used to distribute compiled Java class files and resource files for an application.

Jini A network technology from Sun that allows devices to spontaneously join networks and make their services available to other devices. Although there are some superficial similarities between Jini and JXTA, Jini is heavily dependent on the Java language; JXTA is designed to enable a developer to create a JXTA-compatible application independent of a particular operating system or programming language.

JRE Java Runtime Environment, an environment used to run Java applications. The JRE consists of both a Java Virtual Machine and the standard Java runtime libraries.

JVM Java Virtual Machine, a simulated computer environment used to run Java applications. Java is different from traditional computer languages in that source code is compiled into byte code rather than machine code for a particular processor. A JVM provides a mechanism for translating Java byte code into native machine code at execution time, allowing code to run on any platform that has a JVM implementation.

LAN Local area network, a data-communications network connecting a set of computers in a limited local geographic area.

MD5 Message Digest 5, a hashing algorithm developed by Dr. Ronald Rivest of RSA Security. Hashing algorithms are mathematical one-way functions that convert a stream

of bytes into a unique set of bytes, called a *message digest*. An important property of a good hashing algorithm is it makes it extremely difficult to construct two streams of bytes that result in the same message digest. Message digests are a fixed length, which makes them suitable for uniquely identifying a set of bytes or providing an integrity checksum.

metadata Data that describes other data. In XML, the tags used to mark up data provide metadata that describes the type of data contained by the tag. Metadata provides a higher level representation of information and provides context for the data it describes or contains.

MIME Multipurpose Internet Mail Extensions, a mechanism to allow the exchange of non-ASCII data via Internet mail. MIME defines not only a format for message data, but also MIME types that identify the type of data contained in a MIME message. MIME types are now used by a variety of applications besides email, including HTTP, to allow applications to identify and handle different types of data.

MP3 MPEG-1 Audio Layer 3, an audio compression format defined by the Motion Picture Experts Group.

NAT Network Address Translation, a translation scheme used to protect private internal networks from unauthorized incoming

connections. NAT, usually implemented by a private network's router, translates a private internal IP address to an external public IP address and stores the mapping in the router's translation table.

Incoming connections undergo an inverse mapping procedure; if no mapping exists, the connection is blocked from entering the internal network.

P2P Peer-to-peer, a networking paradigm that enables intermittently connected devices, usually separated from the public network by a firewall, to offer resources to and consume resources from other devices on the network using a common set of communication protocols.

packet A unit of data used in network communications. Typically, a message being sent from one computer to another using IP is divided into a set of packets that are sent across the network in an independent manner to a destination where the message is reassembled.

PDA Personal digital assistant, a handheld computing device that provides some combination of personal information management and communication/networking functionality. Examples of popular PDAs include the Palm Pilot, Handspring Visor, Samsung Yopy, and Compaq iPaq.

peer An entity on the P2P network used to provide access to the resources of the node and consume

resources from other entities on the network.

peer group Peers on a P2P network that join together to serve a common purpose. Peer groups allow peers to segment the network space by application, security, and monitoring requirements.

pipe A virtual communications channel that connects a source endpoint to one or more destination endpoints to permit message exchange.

rendezvous peer A peer that provides simple peers with a way of discovering other peers and advertisements on the P2P network. A rendezvous peer also provides simple peers with the capability to propagate messages within a group, across boundaries between public and private networks. Some rendezvous peers also cache advertisements to reduce network traffic and improve efficiency.

router peer A peer providing routing services to enable peers inside private internal networks behind firewall and NAT equipment to participate in a P2P network.

service A mechanism for providing access to a resource over a network to other peers on a P2P network.

SGML Standard Generalized Markup Language, a method of defining a document language that can be used to mark up documents with metadata using a set of tags. HTML is an SGML-based

document language that defines a set of tags used to mark up documents for presenting within a web browser.

simple peer The simplest type of peer on a P2P network. A simple peer provides resource to and consumes resources from other peers on the network. A simple peer is not responsible for forwarding messages on behalf of other peers or providing third-party information to the network.

SMTP Simple Mail Transport Protocol, a protocol used for exchanging email between servers.

TCP Transport Control Protocol, a protocol that defines rules to guarantee that the packets that form a message arrive at the destination in a timely fashion and are reassembled correctly. TCP is used in conjunction with IP, which provides the function of communicating packets across the network from a source to a destination, in a form called TCP/IP.

TTL Time To Live, a property used to limit the propagation of messages between rendezvous peers. The TTL property of a message defines the maximum number of times that a message should be propagated to other peers. When a rendezvous peer receives a message, it decrements the message's TTL value. If the result is 0, the message is not propagated to other peers; otherwise, the message is propagated using the new TTL value.

UDP User Datagram Packet, a protocol that defines port numbers used to distinguish communication layered on top of IP and checksums to verify data integrity. UDP is used in conjunction with IP, which provides the actual function of communicating packets across the network from a source to a destination. Like IP, UDP is an unreliable protocol that does not guarantee delivery.

UTF-8 Unicode Transformation Format, an encoding scheme used to represent Unicode strings that is specifically optimized for representing ASCII characters.

WSDL Web Services Description Language, an XML-based language used to define the services offered by a server and how to engage those services.

XML eXtensible Markup Language, a language used to define a set of tags that can be used to structure textual data. Tags defined by an XML Document Type Definition are used to mark up data, thereby providing additional information, or metadata, about the data. XML is a simple text-based language that can be easily transformed into other formats using eXtensible Stylesheet Language Transform (XSLT).