

O'REILLY®




ZooKeeper

DISTRIBUTED PROCESS COORDINATION

Flavio Junqueira & Benjamin Reed

Change the world with data.
We'll show you how.
strataconf.com



Strata
CONFERENCE
+
HADOOP
WORLD™

Oct 28 – 30, 2013
New York, NY

O'REILLY®

Strata
CONFERENCE
Making Data Work

Nov 11 – 13, 2013
London, England



O'REILLY®

Strata
CONFERENCE
Making Data Work

Feb 11 – 13, 2014
Santa Clara, CA

O'REILLY®

Strata_{Rx}
CONFERENCE
Data Makes a Difference

April 23–25, 2014
Boston, MA



O'REILLY®

Spreading the knowledge of innovators.

©2013 O'Reilly Media, Inc. O'Reilly logo is a registered trademark of O'Reilly Media, Inc. 13826

ZooKeeper

Flavio Junqueira

Benjamin Reed



Beijing • Cambridge • Farnham • Köln • Sebastopol • Tokyo

Special Upgrade Offer

If you purchased this ebook directly from oreilly.com, you have the following benefits:

- DRM-free ebooks—use your ebooks across devices without restrictions or limitations
- Multiple formats—use on your laptop, tablet, or phone
- Lifetime access, with free updates
- Dropbox syncing—your files, anywhere

If you purchased this ebook from another retailer, you can upgrade your ebook to take advantage of all these benefits for just \$4.99. [Click here](#) to access your ebook upgrade.

Please note that upgrade offers are not available from sample content.

Preface

Building distributed systems is hard. A lot of the applications people use daily, however, depend on such systems, and it doesn't look like we will stop relying on distributed computer systems any time soon. Apache ZooKeeper has been designed to mitigate the task of building robust distributed systems. It has been built around core distributed computing concepts, with its main goal to present the developer with an interface that is simple to understand and program against, thus simplifying the task of building such systems.

Even with ZooKeeper, the task is not trivial—which leads us to this book. This book will get you up to speed on building distributed systems with Apache ZooKeeper. We start with basic concepts that will quickly make you feel like you're a distributed systems expert. Perhaps it will be a bit disappointing to see that it is not that simple when we discuss a bunch of caveats that you need to be aware of. But don't worry; if you understand well the key issues we expose, you'll be on the right track to building great distributed applications.

Audience

This book is aimed at developers of distributed systems and administrators of applications using ZooKeeper in production. We assume knowledge of Java, and try to give you enough background in the principles of distributed systems to use ZooKeeper robustly.

Contents of This Book

Part I covers some motivations for a system like Apache ZooKeeper, and some of the necessary background in distributed systems that you need to use it.

- **Chapter 1, *Introduction***, explains what ZooKeeper can accomplish and how its design supports its mission.
- **Chapter 2, *Getting to Grips with ZooKeeper***, goes over the basic concepts and building blocks. It explains how to get a more concrete idea of what ZooKeeper can do by using the command line.

Part II covers the library calls and programming techniques that programmers need to know. It is useful but not required reading for system administrators. This part focuses on the Java API because it is the most popular. If you are using a different language,

you can read this part to learn the basic techniques and functions, then implement them in a different language. We have an additional chapter covering the C binding for the developers of applications in this language.

- *Chapter 3, Getting Started with the ZooKeeper API*, introduces the Java API.
- *Chapter 4, Dealing with State Change*, explains how to track and react to changes to the state of ZooKeeper.
- *Chapter 5, Dealing with Failure*, shows how to recover from system or network problems.
- *Chapter 6, ZooKeeper Caveat Emptor*, describes some miscellaneous but important considerations you should look for to avoid problems.
- *Chapter 7, The C Client*, introduces the C API, which is the basis for all the non-Java implementations of the ZooKeeper API. Therefore, it's valuable for programmers using any language besides Java.
- *Chapter 8, Curator: A High-Level API for ZooKeeper*, describes a popular high-level interfaces to ZooKeeper.

Part III covers ZooKeeper for system administrators. Programmers might also find it useful, in particular the chapter about internals.

- *Chapter 9, ZooKeeper Internals*, describes some of the choices made by ZooKeeper developers that have an impact on administration tasks.
- *Chapter 10, Running ZooKeeper*, shows how to configure ZooKeeper.

Conventions Used in this Book

The following typographical conventions are used in this book:

Italic

Used for emphasis, new terms, URLs, commands and utilities, and file and directory names.

Constant width

Indicates variables, functions, types, parameters, objects, and other programming constructs.

Constant width bold

Shows commands or other text that should be typed literally by the user. Also used for emphasis in command output.

Constant width italic

Indicates placeholders in code or commands that should be replaced by appropriate values.

TIP

This icon signifies a tip, suggestion, or a general note.

Using Code Examples

Supplemental material (code examples, exercises, etc.) is available for download at <http://bit.ly/zookeeper-code>.

This book is here to help you get your job done. In general, if example code is offered with this book, you may use it in your programs and documentation. You do not need to contact us for permission unless you're reproducing a significant portion of the code. For example, writing a program that uses several chunks of code from this book does not require permission. Selling or distributing a CD-ROM of examples from O'Reilly books does require permission. Answering a question by citing this book and quoting example code does not require permission. Incorporating a significant amount of example code from this book into your product's documentation does require permission.

We appreciate, but do not require, attribution. An attribution usually includes the title, author, publisher, and ISBN. For example: “*ZooKeeper* by Flavio Junqueira and Benjamin Reed (O'Reilly). Copyright 2014 Flavio Junqueira and Benjamin Reed, 978-1-449-36130-3.”

If you feel your use of code examples falls outside fair use or the permission given above, feel free to contact us at permissions@oreilly.com.

Safari® Books Online

Safari Books Online is an on-demand digital library that delivers expert **content** in both book and video form from the world's leading authors in technology and business.

Technology professionals, software developers, web designers, and business and creative professionals use Safari Books Online as their primary resource for research, problem solving, learning, and certification training.

Safari Books Online offers a range of **product mixes** and pricing programs for **organizations**, **government agencies**, and **individuals**. Subscribers have access to thousands of books, training videos, and prepublication manuscripts in one fully

searchable database from publishers like O'Reilly Media, Prentice Hall Professional, Addison-Wesley Professional, Microsoft Press, Sams, Que, Peachpit Press, Focal Press, Cisco Press, John Wiley & Sons, Syngress, Morgan Kaufmann, IBM Redbooks, Packt, Adobe Press, FT Press, Apress, Manning, New Riders, McGraw-Hill, Jones & Bartlett, Course Technology, and dozens **more**. For more information about Safari Books Online, please visit us **online**.

How to Contact Us

Please address comments and questions concerning this book to the publisher:

O'Reilly Media, Inc.
1005 Gravenstein Highway North
Sebastopol, CA 95472
800-998-9938 (in the United States or Canada)
707-829-0515 (international or local)
707-829-0104 (fax)

We have a web page for this book, where we list errata, examples, and any additional information. You can access this page at <http://oreil.ly/zookeeper-orm>.

To comment or ask technical questions about this book, send email to bookquestions@oreilly.com.

For more information about our books, courses, conferences, and news, see our website at <http://www.oreilly.com>.

Find us on Facebook: <http://facebook.com/oreilly>

Follow us on Twitter: <http://twitter.com/oreillymedia>

Watch us on YouTube: <http://www.youtube.com/oreillymedia>

Acknowledgments

We would like to thank our editors, initially Nathan Jepson and later Andy Oram, for the fantastic job they did of getting us to produce this book.

We would like to thank our families and employers for understanding the importance of spending so many hours with this book. We hope you appreciate the outcome.

We would like to thank our reviewers for spending time to give us great comments that helped us to improve the material in this book. They are: Patrick Hunt, Jordan Zimmerman, Donald Miner, Henry Robinson, Isabel Drost-Fromm, and Thawan Kooburat.

ZooKeeper is the collective work of the Apache ZooKeeper community. We work with some really excellent committers and other contributors; it's a privilege to work with

you all. We also want to give a big thanks to all of the ZooKeeper users who have reported bugs and given us so much feedback and encouragement over the years.