



# Apache HBase Primer

---

Deepak Vohra

Apress®

# Apache HBase Primer



Deepak Vohra

Apress®

## ***Apache HBase Primer***

Deepak Vohra  
White Rock, British Columbia  
Canada

ISBN-13 (pbk): 978-1-4842-2423-6  
DOI 10.1007/978-1-4842-2424-3

ISBN-13 (electronic): 978-1-4842-2424-3

Library of Congress Control Number: 2016959189

Copyright © 2016 by Deepak Vohra

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director: Welmoed Spahr

Lead Editor: Steve Anglin

Technical Reviewer: Massimo Nardone

Editorial Board: Steve Anglin, Pramila Balan, Laura Berendson, Aaron Black,

Louise Corrigan, Jonathan Gennick, Robert Hutchinson, Celestin Suresh John,

Nikhil Karkal, James Markham, Susan McDermott, Matthew Moodie, Natalie Pao,

Gwenan Spearing

Coordinating Editor: Mark Powers

Copy Editor: Mary Behr

Compositor: SPi Global

Indexer: SPi Global

Artist: SPi Global

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail [orders-ny@springer-sbm.com](mailto:orders-ny@springer-sbm.com), or visit [www.springeronline.com](http://www.springeronline.com). Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a Delaware corporation.

For information on translations, please e-mail [rights@apress.com](mailto:rights@apress.com), or visit [www.apress.com](http://www.apress.com).

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles.

For more information, reference our Special Bulk Sales–eBook Licensing web page at [www.apress.com/bulk-sales](http://www.apress.com/bulk-sales).

Any source code or other supplementary materials referenced by the author in this text are available to readers at [www.apress.com](http://www.apress.com). For detailed information about how to locate your book's source code, go to [www.apress.com/source-code/](http://www.apress.com/source-code/). Readers can also access source code at SpringerLink in the Supplementary Material section for each chapter.

Printed on acid-free paper

# Contents at a Glance

<b>About the Author .....</b>	<b>xiii</b>
<b>About the Technical Reviewer .....</b>	<b>xv</b>
<b>Introduction .....</b>	<b>xvii</b>
<b>■ Part I: Core Concepts .....</b>	<b>1</b>
<b>■ Chapter 1: Fundamental Characteristics .....</b>	<b>3</b>
<b>■ Chapter 2: Apache HBase and HDFS .....</b>	<b>9</b>
<b>■ Chapter 3: Application Characteristics.....</b>	<b>45</b>
<b>■ Part II: Data Model.....</b>	<b>49</b>
<b>■ Chapter 4: Physical Storage .....</b>	<b>51</b>
<b>■ Chapter 5: Column Family and Column Qualifier.....</b>	<b>53</b>
<b>■ Chapter 6: Row Versioning .....</b>	<b>59</b>
<b>■ Chapter 7: Logical Storage .....</b>	<b>63</b>
<b>■ Part III: Architecture .....</b>	<b>67</b>
<b>■ Chapter 8: Major Components of a Cluster.....</b>	<b>69</b>
<b>■ Chapter 9: Regions .....</b>	<b>75</b>
<b>■ Chapter 10: Finding a Row in a Table .....</b>	<b>81</b>
<b>■ Chapter 11: Compactions .....</b>	<b>87</b>
<b>■ Chapter 12: Region Failover .....</b>	<b>99</b>
<b>■ Chapter 13: Creating a Column Family .....</b>	<b>105</b>

- **Part IV: Schema Design** ..... 109
- **Chapter 14: Region Splitting**..... 111
- **Chapter 15: Defining the Row Keys** ..... 117
- **Part V: Apache HBase Java API** ..... 121
- **Chapter 16: The HBaseAdmin Class**..... 123
- **Chapter 17: Using the Get Class** ..... 129
- **Chapter 18: Using the HTable Class** ..... 133
- **Part VI: Administration** ..... 135
- **Chapter 19: Using the HBase Shell** ..... 137
- **Chapter 20: Bulk Loading Data** ..... 145
- Index**..... 149

# Contents

<b>About the Author .....</b>	<b>xiii</b>
<b>About the Technical Reviewer .....</b>	<b>xv</b>
<b>Introduction .....</b>	<b>xvii</b>
<b>■ Part I: Core Concepts .....</b>	<b>1</b>
<b>■ Chapter 1: Fundamental Characteristics .....</b>	<b>3</b>
Distributed.....	3
Big Data Store .....	3
Non-Relational.....	3
Flexible Data Model.....	4
Scalable.....	4
Roles in Hadoop Big Data Ecosystem.....	5
How Is Apache HBase Different from a Traditional RDBMS? .....	5
Summary.....	8
<b>■ Chapter 2: Apache HBase and HDFS .....</b>	<b>9</b>
Overview .....	9
Storing Data .....	14
HFile Data files- HFile v1 .....	15
HBase Blocks .....	17
Key Value Format .....	18
HFile v2 .....	19
Encoding.....	20

- Compaction ..... 21
- KeyValue Class ..... 21
- Data Locality..... 24
- Table Format ..... 25
- HBase Ecosystem..... 25
- HBase Services ..... 26
- Auto-sharding..... 27
- The Write Path to Create a Table ..... 27
- The Write Path to Insert Data ..... 28
- The Write Path to Append-Only R/W ..... 29
- The Read Path for Reading Data ..... 30
- The Read Path Append-Only to Random R/W..... 30
- HFile Format..... 30
- Data Block Encoding ..... 31
- Compactions ..... 32
- Snapshots ..... 32
- The HFileSystem Class..... 33
- Scaling ..... 33
- HBase Java Client API..... 35
- Random Access..... 36
- Data Files (HFile) ..... 36
- Reference Files/Links..... 37
- Write-Ahead Logs..... 38
- Data Locality..... 38
- Checksums..... 40
- Data Locality for HBase ..... 42

MemStore.....	42
Summary.....	43
■ <b>Chapter 3: Application Characteristics.....</b>	<b>45</b>
Summary.....	47
■ <b>Part II: Data Model.....</b>	<b>49</b>
■ <b>Chapter 4: Physical Storage .....</b>	<b>51</b>
Summary.....	52
■ <b>Chapter 5: Column Family and Column Qualifier .....</b>	<b>53</b>
Summary.....	57
■ <b>Chapter 6: Row Versioning .....</b>	<b>59</b>
Versions Sorting .....	61
Summary.....	62
■ <b>Chapter 7: Logical Storage .....</b>	<b>63</b>
Summary.....	65
■ <b>Part III: Architecture .....</b>	<b>67</b>
■ <b>Chapter 8: Major Components of a Cluster.....</b>	<b>69</b>
Master .....	70
RegionServers .....	70
ZooKeeper .....	71
Regions .....	72
Write-Ahead Log.....	72
Store.....	72
HDFS.....	73
Clients .....	73
Summary.....	73



- **Chapter 9: Regions** ..... **75**
  - How Many Regions?..... 76
  - Compactions ..... 76
  - Region Assignment..... 76
  - Failover..... 77
  - Region Locality ..... 77
  - Distributed Datastore ..... 77
  - Partitioning ..... 77
  - Auto Sharding and Scalability ..... 78
  - Region Splitting ..... 78
  - Manual Splitting ..... 79
  - Pre-Splitting ..... 79
  - Load Balancing..... 79
  - Preventing Hotspots ..... 80
  - Summary ..... 80
- **Chapter 10: Finding a Row in a Table** ..... **81**
  - Block Cache..... 82
  - The hbase:meta Table ..... 83
  - Summary ..... 85
- **Chapter 11: Compactions** ..... **87**
  - Minor Compactions ..... 87
  - Major Compactions ..... 88
  - Compaction Policy..... 88
  - Function and Purpose ..... 89
  - Versions and Compactions ..... 90
  - Delete Markers and Compactions ..... 90
  - Expired Rows and Compactions ..... 90

Region Splitting and Compactions .....	90
Number of Regions and Compactions .....	91
Data Locality and Compactions .....	91
Write Throughput and Compactions .....	91
Encryption and Compactions.....	91
Configuration Properties .....	92
Summary .....	97
<b>■ Chapter 12: Region Failover .....</b>	<b>99</b>
The Role of the ZooKeeper .....	99
HBase Resilience.....	99
Phases of Failover .....	100
Failure Detection .....	102
Data Recovery .....	102
Regions Reassignment.....	103
Failover and Data Locality .....	103
Configuration Properties .....	103
Summary .....	103
<b>■ Chapter 13: Creating a Column Family .....</b>	<b>105</b>
Cardinality .....	105
Number of Column Families .....	106
Column Family Compression.....	106
Column Family Block Size .....	106
Bloom Filters .....	106
IN_MEMORY .....	107
MAX_LENGTH and MAX_VERSIONS.....	107
Summary .....	107

- **Part IV: Schema Design ..... 109**
- **Chapter 14: Region Splitting..... 111**
  - Managed Splitting ..... 112
  - Pre-Splitting ..... 113
  - Configuration Properties ..... 113
  - Summary ..... 116
- **Chapter 15: Defining the Row Keys ..... 117**
  - Table Key Design ..... 117
  - Filters ..... 118
    - FirstKeyOnlyFilter Filter ..... 118
    - KeyOnlyFilter Filter ..... 118
  - Bloom Filters ..... 118
  - Scan Time..... 118
  - Sequential Keys..... 118
  - Defining the Row Keys for Locality ..... 119
  - Summary ..... 119
- **Part V: Apache HBase Java API ..... 121**
- **Chapter 16: The HBaseAdmin Class..... 123**
  - Summary ..... 127
- **Chapter 17: Using the Get Class ..... 129**
  - Summary ..... 132
- **Chapter 18: Using the HTable Class ..... 133**
  - Summary ..... 134

■ **Part VI: Administration** ..... **135**

■ **Chapter 19: Using the HBase Shell** ..... **137**

    Creating a Table..... 137

    Altering a Table..... 138

    Adding Table Data..... 139

    Describing a Table ..... 139

    Finding If a Table Exists..... 139

    Listing Tables..... 139

    Scanning a Table ..... 140

    Enabling and Disabling a Table..... 141

    Dropping a Table..... 141

    Counting the Number of Rows in a Table ..... 141

    Getting Table Data ..... 141

    Truncating a Table ..... 142

    Deleting Table Data ..... 142

    Summary..... 143

■ **Chapter 20: Bulk Loading Data** ..... **145**

    Summary..... 147

**Index**..... **149**

# About the Author



**Deepak Vohra** is a consultant and a principal member of the NuBean software company. Deepak is a Sun-certified Java programmer and Web component developer. He has worked in the fields of XML, Java programming, and Java EE for over seven years. Deepak is the coauthor of *Pro XML Development with Java Technology* (Apress, 2006). Deepak is also the author of the *JDBC 4.0* and *Oracle JDeveloper for J2EE Development, Processing XML Documents with Oracle JDeveloper 11g, EJB 3.0 Database Persistence with Oracle Fusion Middleware 11g, and Java EE Development in Eclipse IDE* (Packt Publishing). He also served as the technical reviewer on *WebLogic: The Definitive Guide* (O'Reilly Media, 2004) and *Ruby Programming for the Absolute Beginner* (Cengage Learning PTR, 2007).

# About the Technical Reviewer



**Massimo Nardone** has more than 22 years of experience in security, web/mobile development, and cloud and IT architecture. His true IT passions are security and Android. He has been programming and teaching how to program with Android, Perl, PHP, Java, VB, Python, C/C++, and MySQL for more than 20 years. Technical skills include security, Android, cloud, Java, MySQL, Drupal, Cobol, Perl, web and mobile development, MongoDB, D3, Joomla, Couchbase, C/C++, WebGL, Python, Pro Rails, Django CMS, Jekyll, Scratch, etc.

He currently works as Chief Information Security Office (CISO) for Cargotec Oyj. He holds four international patents (PKI, SIP, SAML, and Proxy areas).

He worked as a visiting lecturer and supervisor for exercises at the Networking Laboratory of the Helsinki University of Technology (Aalto University). He has also worked as a Project Manager, Software Engineer, Research Engineer, Chief Security Architect, Information Security Manager, PCI/SCADA Auditor, and Senior Lead IT Security/Cloud/SCADA Architect for many years. He holds a Master of Science degree in Computing Science from the University of Salerno, Italy.

Massimo has reviewed more than 40 IT books for different publishing companies, and he is the coauthor of *Pro Android Games* (Apress, 2015).

# Introduction

Apache HBase is an open source NoSQL database based on the wide-column data store model. HBase was initially released in 2008. While many NoSQL databases are available, Apache HBase is the database for the Apache Hadoop ecosystem.

HBase supports most of the commonly used programming languages such as C, C++, PHP, and Java. The implementation language of HBase is Java. HBase provides access support with Java API, RESTful HTTP API, and Thrift.

Some of the other Apache HBase books have a practical orientation and do not discuss HBase concepts in much detail. In this primer level book, I shall discuss Apache HBase concepts. For practical use of Apache HBase, refer another Apress book: *Practical Hadoop Ecosystem*.