

# Hadoop

## IN PRACTICE

SECOND EDITION

Alex Holmes



INCLUDES 104 TECHNIQUES

 MANNING

## Praise for the First Edition of *Hadoop in Practice*

*A new book from Manning, Hadoop in Practice, is definitely the most modern book on the topic. Important subjects, like what commercial variants such as MapR offer, and the many different releases and APIs get uniquely good coverage in this book.*

—Ted Dunning, Chief Application Architect, MapR Technologies

*Comprehensive coverage of advanced Hadoop usage, including high-quality code samples.*

—Chris Nauroth, Senior Staff Software Engineer  
The Walt Disney Company

*A very pragmatic and broad overview of Hadoop and the Hadoop tools ecosystem, with a wide set of interesting topics that tickle the creative brain.*

—Mark Kemna, Chief Technology Officer, Brilig

*A practical introduction to the Hadoop ecosystem.*

—Philipp K. Janert, Principal Value, LLC

*This book is the horizontal roof that each of the pillars of individual Hadoop technology books hold. It expertly ties together all the Hadoop ecosystem technologies.*

—Ayon Sinha, Big Data Architect, Britely

*I would take this book on my path to the future.*

—Alexey Gayduk, Senior Software Engineer, Grid Dynamics

*A high-quality and well-written book that is packed with useful examples. The breadth and detail of the material is by far superior to any other Hadoop reference guide. It is perfect for anyone who likes to learn new tools/technologies while following pragmatic, real-world examples.*

—Amazon reviewer



*Hadoop in Practice*  
*Second Edition*

ALEX HOLMES



MANNING  
Shelter Island

For online information and ordering of this and other Manning books, please visit [www.manning.com](http://www.manning.com). The publisher offers discounts on this book when ordered in quantity. For more information, please contact

Special Sales Department  
Manning Publications Co.  
20 Baldwin Road  
PO Box 761  
Shelter Island, NY 11964  
Email: [orders@manning.com](mailto:orders@manning.com)

©2015 by Manning Publications Co. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by means electronic, mechanical, photocopying, or otherwise, without prior written permission of the publisher.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in the book, and Manning Publications was aware of a trademark claim, the designations have been printed in initial caps or all caps.

© Recognizing the importance of preserving what has been written, it is Manning's policy to have the books we publish printed on acid-free paper, and we exert our best efforts to that end. Recognizing also our responsibility to conserve the resources of our planet, Manning books are printed on paper that is at least 15 percent recycled and processed without the use of elemental chlorine.



Manning Publications Co.  
20 Baldwin Road  
Shelter Island, NY 11964

Development editor: Cynthia Kane  
Copyeditor: Andy Carroll  
Proofreader: Melody Dolab  
Typesetter: Gordan Salinovic  
Cover designer: Marija Tudor

ISBN 9781617292224

Printed in the United States of America

1 2 3 4 5 6 7 8 9 10 – EBM – 19 18 17 16 15 14

# *brief contents*

---

<b>PART 1</b>	<b>BACKGROUND AND FUNDAMENTALS .....</b>	<b>1</b>
	1 ■ Hadoop in a heartbeat	3
	2 ■ Introduction to YARN	22
<b>PART 2</b>	<b>DATA LOGISTICS .....</b>	<b>59</b>
	3 ■ Data serialization—working with text and beyond	61
	4 ■ Organizing and optimizing data in HDFS	139
	5 ■ Moving data into and out of Hadoop	174
<b>PART 3</b>	<b>BIG DATA PATTERNS .....</b>	<b>253</b>
	6 ■ Applying MapReduce patterns to big data	255
	7 ■ Utilizing data structures and algorithms at scale	302
	8 ■ Tuning, debugging, and testing	337
<b>PART 4</b>	<b>BEYOND MAPREDUCE .....</b>	<b>385</b>
	9 ■ SQL on Hadoop	387
	10 ■ Writing a YARN application	425



# contents

---

*preface* xv  
*acknowledgments* xvii  
*about this book* xviii  
*about the cover illustration* xxiii

## **PART 1 BACKGROUND AND FUNDAMENTALS .....1**

### **1 Hadoop in a heartbeat 3**

#### 1.1 What is Hadoop? 4

*Core Hadoop components* 5 ▪ *The Hadoop ecosystem* 10  
*Hardware requirements* 11 ▪ *Hadoop distributions* 12 ▪ *Who's using Hadoop?* 14 ▪ *Hadoop limitations* 15

#### 1.2 Getting your hands dirty with MapReduce 17

#### 1.3 Summary 21

### **2 Introduction to YARN 22**

#### 2.1 YARN overview 23

*Why YARN?* 24 ▪ *YARN concepts and components* 26  
*YARN configuration* 29

TECHNIQUE 1 Determining the configuration of your cluster 29

*Interacting with YARN* 31



TECHNIQUE	2	Running a command on your YARN cluster	31
TECHNIQUE	3	Accessing container logs	32
TECHNIQUE	4	Aggregating container log files	36
		<i>YARN challenges</i>	39
2.2		YARN and MapReduce	40
		<i>Dissecting a YARN MapReduce application</i>	40
		▪ <i>Configuration</i>	42
		<i>Backward compatibility</i>	46
TECHNIQUE	5	Writing code that works on Hadoop versions 1 and 2	47
		<i>Running a job</i>	48
TECHNIQUE	6	Using the command line to run a job	49
		<i>Monitoring running jobs and viewing archived jobs</i>	49
		<i>Uber jobs</i>	50
TECHNIQUE	7	Running small MapReduce jobs	50
2.3		YARN applications	52
		<i>NoSQL</i>	53
		▪ <i>Interactive SQL</i>	54
		▪ <i>Graph processing</i>	54
		<i>Real-time data processing</i>	55
		▪ <i>Bulk synchronous parallel</i>	55
		<i>MPI</i>	56
		▪ <i>In-memory</i>	56
		▪ <i>DAG execution</i>	56
2.4		Summary	57

## PART 2 DATA LOGISTICS.....59

### 3 Data serialization—working with text and beyond 61

3.1		Understanding inputs and outputs in MapReduce	62
		<i>Data input</i>	63
		▪ <i>Data output</i>	66
3.2		Processing common serialization formats	68
		<i>XML</i>	69
TECHNIQUE	8	MapReduce and XML	69
		<i>JSON</i>	72
TECHNIQUE	9	MapReduce and JSON	73
3.3		Big data serialization formats	76
		<i>Comparing SequenceFile, Protocol Buffers, Thrift, and Avro</i>	76
		<i>SequenceFile</i>	78
TECHNIQUE	10	Working with SequenceFiles	80
TECHNIQUE	11	Using SequenceFiles to encode Protocol Buffers	87
		<i>Protocol Buffers</i>	91
		▪ <i>Thrift</i>	92
		▪ <i>Avro</i>	93
TECHNIQUE	12	Avro's schema and code generation	93

TECHNIQUE	13	Selecting the appropriate way to use Avro in MapReduce	98
TECHNIQUE	14	Mixing Avro and non-Avro data in MapReduce	99
TECHNIQUE	15	Using Avro records in MapReduce	102
TECHNIQUE	16	Using Avro key/value pairs in MapReduce	104
TECHNIQUE	17	Controlling how sorting works in MapReduce	108
TECHNIQUE	18	Avro and Hive	108
TECHNIQUE	19	Avro and Pig	111
3.4		Columnar storage	113
		<i>Understanding object models and storage formats</i>	115
		▪ <i>Parquet and the Hadoop ecosystem</i>	116
		▪ <i>Parquet block and page sizes</i>	117
TECHNIQUE	20	Reading Parquet files via the command line	117
TECHNIQUE	21	Reading and writing Avro data in Parquet with Java	119
TECHNIQUE	22	Parquet and MapReduce	120
TECHNIQUE	23	Parquet and Hive/Impala	125
TECHNIQUE	24	Pushdown predicates and projection with Parquet	126
		<i>Parquet limitations</i>	128
3.5		Custom file formats	129
		<i>Input and output formats</i>	129
TECHNIQUE	25	Writing input and output formats for CSV	129
		<i>The importance of output committing</i>	137
3.6		Chapter summary	138

## 4 Organizing and optimizing data in HDFS 139

4.1		Data organization	140
		<i>Directory and file layout</i>	140
		▪ <i>Data tiers</i>	141
		▪ <i>Partitioning</i>	142
TECHNIQUE	26	Using MultipleOutputs to partition your data	142
TECHNIQUE	27	Using a custom MapReduce partitioner	145
		<i>Compacting</i>	148
TECHNIQUE	28	Using filecrush to compact data	149
TECHNIQUE	29	Using Avro to store multiple small binary files	151
		<i>Atomic data movement</i>	157
4.2		Efficient storage with compression	158
TECHNIQUE	30	Picking the right compression codec for your data	159

TECHNIQUE	31	Compression with HDFS, MapReduce, Pig, and Hive	163
TECHNIQUE	32	Splittable LZOP with MapReduce, Hive, and Pig	168
4.3		Chapter summary	173

## 5 *Moving data into and out of Hadoop* 174

5.1		Key elements of data movement	175
5.2		Moving data into Hadoop	177
		<i>Roll your own ingest</i>	177
TECHNIQUE	33	Using the CLI to load files	178
TECHNIQUE	34	Using REST to load files	180
TECHNIQUE	35	Accessing HDFS from behind a firewall	183
TECHNIQUE	36	Mounting Hadoop with NFS	186
TECHNIQUE	37	Using DistCp to copy data within and between clusters	188
TECHNIQUE	38	Using Java to load files	194
		<i>Continuous movement of log and binary files into HDFS</i>	196
TECHNIQUE	39	Pushing system log messages into HDFS with Flume	197
TECHNIQUE	40	An automated mechanism to copy files into HDFS	204
TECHNIQUE	41	Scheduling regular ingress activities with Oozie	209
		<i>Databases</i>	214
TECHNIQUE	42	Using Sqoop to import data from MySQL	215
		<i>HBase</i>	227
TECHNIQUE	43	HBase ingress into HDFS	227
TECHNIQUE	44	MapReduce with HBase as a data source	230
		<i>Importing data from Kafka</i>	232
5.3		Moving data into Hadoop	234
TECHNIQUE	45	Using Camus to copy Avro data from Kafka into HDFS	234
5.4		Moving data out of Hadoop	241
		<i>Roll your own egress</i>	241
TECHNIQUE	46	Using the CLI to extract files	241
TECHNIQUE	47	Using REST to extract files	242
TECHNIQUE	48	Reading from HDFS when behind a firewall	243
TECHNIQUE	49	Mounting Hadoop with NFS	243
TECHNIQUE	50	Using DistCp to copy data out of Hadoop	244

TECHNIQUE	51	Using Java to extract files	245
		<i>Automated file egress</i>	246
TECHNIQUE	52	An automated mechanism to export files from HDFS	246
		<i>Databases</i>	247
TECHNIQUE	53	Using Sqoop to export data to MySQL	247
		<i>NoSQL</i>	251
5.5		Chapter summary	252

## PART 3 BIG DATA PATTERNS.....253

### 6 *Applying MapReduce patterns to big data* 255

6.1		Joining	256
TECHNIQUE	54	Picking the best join strategy for your data	257
TECHNIQUE	55	Filters, projections, and pushdowns	259
		<i>Map-side joins</i>	260
TECHNIQUE	56	Joining data where one dataset can fit into memory	261
TECHNIQUE	57	Performing a semi-join on large datasets	264
TECHNIQUE	58	Joining on presorted and prepartitioned data	269
		<i>Reduce-side joins</i>	271
TECHNIQUE	59	A basic repartition join	271
TECHNIQUE	60	Optimizing the repartition join	275
TECHNIQUE	61	Using Bloom filters to cut down on shuffled data	279
		<i>Data skew in reduce-side joins</i>	283
TECHNIQUE	62	Joining large datasets with high join-key cardinality	284
TECHNIQUE	63	Handling skews generated by the hash partitioner	286
6.2		Sorting	287
		<i>Secondary sort</i>	288
TECHNIQUE	64	Implementing a secondary sort	289
		<i>Total order sorting</i>	294
TECHNIQUE	65	Sorting keys across multiple reducers	294
6.3		Sampling	297
TECHNIQUE	66	Writing a reservoir-sampling InputFormat	297
6.4		Chapter summary	301

<b>7</b>	<b><i>Utilizing data structures and algorithms at scale</i></b>	<b>302</b>
7.1	Modeling data and solving problems with graphs	303
	<i>Modeling graphs</i>	304
	▪ <i>Shortest-path algorithm</i>	304
TECHNIQUE	67 Find the shortest distance between two users	305
	<i>Friends-of-friends algorithm</i>	313
TECHNIQUE	68 Calculating FoFs	313
	<i>Using Giraph to calculate PageRank over a web graph</i>	319
7.2	Modeling data and solving problems with graphs	321
TECHNIQUE	69 Calculate PageRank over a web graph	322
7.3	Bloom filters	326
TECHNIQUE	70 Parallelized Bloom filter creation in MapReduce	328
7.4	HyperLogLog	333
	<i>A brief introduction to HyperLogLog</i>	333
TECHNIQUE	71 Using HyperLogLog to calculate unique counts	335
7.5	Chapter summary	336
<b>8</b>	<b><i>Tuning, debugging, and testing</i></b>	<b>337</b>
8.1	Measure, measure, measure	338
8.2	Tuning MapReduce	339
	<i>Common inefficiencies in MapReduce jobs</i>	339
TECHNIQUE	72 Viewing job statistics	340
	<i>Map optimizations</i>	343
TECHNIQUE	73 Data locality	343
TECHNIQUE	74 Dealing with a large number of input splits	344
TECHNIQUE	75 Generating input splits in the cluster with YARN	346
	<i>Shuffle optimizations</i>	347
TECHNIQUE	76 Using the combiner	347
TECHNIQUE	77 Blazingly fast sorting with binary comparators	349
TECHNIQUE	78 Tuning the shuffle internals	353
	<i>Reducer optimizations</i>	356
TECHNIQUE	79 Too few or too many reducers	356
	<i>General tuning tips</i>	357

TECHNIQUE	80	Using stack dumps to discover unoptimized user code	358
TECHNIQUE	81	Profiling your map and reduce tasks	360
8.3	Debugging		362
		<i>Accessing container log output</i>	362
TECHNIQUE	82	Examining task logs	362
		<i>Accessing container start scripts</i>	363
TECHNIQUE	83	Figuring out the container startup command	363
		<i>Debugging OutOfMemory errors</i>	365
TECHNIQUE	84	Force container JVMs to generate a heap dump	365
		<i>MapReduce coding guidelines for effective debugging</i>	365
TECHNIQUE	85	Augmenting MapReduce code for better debugging	365
8.4	Testing MapReduce jobs		368
		<i>Essential ingredients for effective unit testing</i>	368 ■ MRUnit 370
TECHNIQUE	86	Using MRUnit to unit-test MapReduce	371
		<i>LocalJobRunner</i>	378
TECHNIQUE	87	Heavyweight job testing with the LocalJobRunner	378
		<i>MiniMRYarnCluster</i>	381
TECHNIQUE	88	Using MiniMRYarnCluster to test your jobs	381
		<i>Integration and QA testing</i>	382
8.5	Chapter summary		383

## PART 4 BEYOND MAPREDUCE .....385

### 9 SQL on Hadoop 387

9.1	Hive		388
		<i>Hive basics</i>	388 ■ <i>Reading and writing data</i>
			391
TECHNIQUE	89	Working with text files	391
TECHNIQUE	90	Exporting data to local disk	395
		<i>User-defined functions in Hive</i>	396
TECHNIQUE	91	Writing UDFs	396
		<i>Hive performance</i>	399
TECHNIQUE	92	Partitioning	399
TECHNIQUE	93	Tuning Hive joins	404

- 9.2 Impala 409
  - Impala vs. Hive* 410 ▪ *Impala basics* 410
  - TECHNIQUE 94 Working with text 410
  - TECHNIQUE 95 Working with Parquet 412
  - TECHNIQUE 96 Refreshing metadata 413
    - User-defined functions in Impala* 414
  - TECHNIQUE 97 Executing Hive UDFs in Impala 415
- 9.3 Spark SQL 416
  - Spark 101* 417 ▪ *Spark on Hadoop* 419 ▪ *SQL with Spark* 419
  - TECHNIQUE 98 Calculating stock averages with Spark SQL 420
  - TECHNIQUE 99 Language-integrated queries 422
  - TECHNIQUE 100 Hive and Spark SQL 423
- 9.4 Chapter summary 423

## 10 *Writing a YARN application* 425

- 10.1 Fundamentals of building a YARN application 426
  - Actors* 426 ▪ *The mechanics of a YARN application* 427
- 10.2 Building a YARN application to collect cluster statistics 429
  - TECHNIQUE 101 A bare-bones YARN client 429
  - TECHNIQUE 102 A bare-bones ApplicationMaster 434
  - TECHNIQUE 103 Running the application and accessing logs 438
  - TECHNIQUE 104 Debugging using an unmanaged application master 440
- 10.3 Additional YARN application capabilities 443
  - RPC between components* 443 ▪ *Service discovery* 444
  - Checkpointing application progress* 444 ▪ *Avoiding split-brain* 444
  - Long-running applications* 444 ▪ *Security* 445
- 10.4 YARN programming abstractions 445
  - Twill* 446 ▪ *Spring* 448 ▪ *REEF* 450 ▪ *Picking a YARN API abstraction* 450
- 10.5 Summary 450
- appendix Installing Hadoop and friends* 451
- index* 475

*bonus chapters available for download from [www.manning.com/holmes2](http://www.manning.com/holmes2)*

*chapter 11 Integrating R and Hadoop for statistics and more*

*chapter 12 Predictive analytics with Mahout*

## *preface*

---

I first encountered Hadoop in the fall of 2008 when I was working on an internet crawl-and-analysis project at Verisign. We were making discoveries similar to those that Doug Cutting and others at Nutch had made several years earlier about how to efficiently store and manage terabytes of crawl-and-analyzed data. At the time, we were getting by with our homegrown distributed system, but the influx of a new data stream and requirements to join that stream with our crawl data couldn't be supported by our existing system in the required timeline.

After some research, we came across the Hadoop project, which seemed to be a perfect fit for our needs—it supported storing large volumes of data and provided a compute mechanism to combine them. Within a few months, we built and deployed a MapReduce application encompassing a number of MapReduce jobs, woven together with our own MapReduce workflow management system, onto a small cluster of 18 nodes. It was a revelation to observe our MapReduce jobs crunching through our data in minutes. Of course, what we weren't expecting was the amount of time that we would spend debugging and performance-tuning our MapReduce jobs. Not to mention the new roles we took on as production administrators—the biggest surprise in this role was the number of disk failures we encountered during those first few months supporting production.

As our experience and comfort level with Hadoop grew, we continued to build more of our functionality using Hadoop to help with our scaling challenges. We also started to evangelize the use of Hadoop within our organization and helped kick-start other projects that were also facing big data challenges.



The greatest challenge we faced when working with Hadoop, and specifically MapReduce, was relearning how to solve problems with it. MapReduce is its own flavor of parallel programming, and it's quite different from the in-JVM programming that we were accustomed to. The first big hurdle was training our brains to think MapReduce, a topic which the book *Hadoop in Action* by Chuck Lam (Manning Publications, 2010) covers well.

After one is used to thinking in MapReduce, the next challenge is typically related to the logistics of working with Hadoop, such as how to move data in and out of HDFS and effective and efficient ways to work with data in Hadoop. These areas of Hadoop haven't received much coverage, and that's what attracted me to the potential of this book—the chance to go beyond the fundamental word-count Hadoop uses and covering some of the trickier and dirtier aspects of Hadoop.

As I'm sure many authors have experienced, I went into this project confidently believing that writing this book was just a matter of transferring my experiences onto paper. Boy, did I get a reality check, but not altogether an unpleasant one, because writing introduced me to new approaches and tools that ultimately helped better my own Hadoop abilities. I hope that you get as much out of reading this book as I did writing it.

## *acknowledgments*

---

First and foremost, I want to thank Michael Noll, who pushed me to write this book. He provided invaluable insights into how to structure the content of the book, reviewed my early chapter drafts, and helped mold the book. I can't express how much his support and encouragement has helped me throughout the process.

I'm also indebted to Cynthia Kane, my development editor at Manning, who coached me through writing this book and provided invaluable feedback on my work. Among the many notable "aha!" moments I had when working with Cynthia, the biggest one was when she steered me into using visual aids to help explain some of the complex concepts in this book.

All of the Manning staff were a pleasure to work with, and a special shout out goes to Troy Mott, Nick Chase, Tara Walsh, Bob Herbstman, Michael Stephens, Marjan Bace, Maureen Spencer, and Kevin Sullivan.

I also want to say a big thank you to all the reviewers of this book: Adam Kawa, Andrea Tarocchi, Anna Lahoud, Arthur Zubarev, Edward Ribeiro, Fillipe Massuda, Gerd Koenig, Jeet Marwah, Leon Portman, Mohamed Diouf, Muthuswamy Manigandan, Rodrigo Abreu, and Serega Sheypack. Jonathan Siedman, the primary technical reviewer, did a great job of reviewing the entire book.

Many thanks to Josh Wills, the creator of Crunch, who kindly looked over the chapter that covered that topic. And more thanks go to Josh Patterson, who reviewed my Mahout chapter.

Finally, a special thanks to my wife, Michal, who had to put up with a cranky husband working crazy hours. She was a source of encouragement throughout the entire process.

## *about this book*

---

Doug Cutting, the creator of Hadoop, likes to call Hadoop the kernel for big data, and I would tend to agree. With its distributed storage and compute capabilities, Hadoop is fundamentally an enabling technology for working with huge datasets. Hadoop provides a bridge between structured (RDBMS) and unstructured (log files, XML, text) data and allows these datasets to be easily joined together. This has evolved from traditional use cases, such as combining OLTP and log files, to more sophisticated uses, such as using Hadoop for data warehousing (exemplified by Facebook) and the field of data science, which studies and makes new discoveries about data.

This book collects a number of intermediary and advanced Hadoop examples and presents them in a problem/solution format. Each technique addresses a specific task you'll face, like using Flume to move log files into Hadoop or using Mahout for predictive analysis. Each problem is explored step by step, and as you work through them, you'll find yourself growing more comfortable with Hadoop and at home in the world of big data.

This hands-on book targets users who have some practical experience with Hadoop and understand the basic concepts of MapReduce and HDFS. Manning's *Hadoop in Action* by Chuck Lam contains the necessary prerequisites to understand and apply the techniques covered in this book.

Many techniques in this book are Java-based, which means readers are expected to possess an intermediate-level knowledge of Java. An excellent text for all levels of Java users is *Effective Java*, Second Edition by Joshua Bloch (Addison-Wesley, 2008).

## **Roadmap**

This book has 10 chapters divided into four parts.

Part 1 contains two chapters that form the introduction to this book. They review Hadoop basics and look at how to get Hadoop up and running on a single host. YARN, which is new in Hadoop version 2, is also examined, and some operational tips are provided for performing basic functions in YARN.

Part 2, “Data logistics,” consists of three chapters that cover the techniques and tools required to deal with data fundamentals, how to work with various data formats, how to organize and optimize your data, and getting data into and out of Hadoop. Picking the right format for your data and determining how to organize data in HDFS are the first items you’ll need to address when working with Hadoop, and they’re covered in chapters 3 and 4 respectively. Getting data into Hadoop is one of the bigger hurdles commonly encountered when working with Hadoop, and chapter 5 is dedicated to looking at a variety of tools that work with common enterprise data sources.

Part 3 is called “Big data patterns,” and it looks at techniques to help you work effectively with large volumes of data. Chapter 6 covers how to represent data such as graphs for use with MapReduce, and it looks at several algorithms that operate on graph data. Chapter 7 looks at more advanced data structures and algorithms such as graph processing and using HyperLogLog for working with large datasets. Chapter 8 looks at how to tune, debug, and test MapReduce performance issues, and it also covers a number of techniques to help make your jobs run faster.

Part 4 is titled “Beyond MapReduce,” and it examines a number of technologies that make it easier to work with Hadoop. Chapter 9 covers the most prevalent and promising SQL technologies for data processing on Hadoop, and Hive, Impala, and Spark SQL are examined. The final chapter looks at how to write your own YARN application, and it provides some insights into some of the more advanced features you can use in your applications.

The appendix covers instructions for the source code that accompanies this book, as well as installation instructions for Hadoop and all the other related technologies covered in the book.

Finally, there are two bonus chapters available from the publisher’s website at [www.manning.com/HadoopinPracticeSecondEdition](http://www.manning.com/HadoopinPracticeSecondEdition): chapter 11 “Integrating R and Hadoop for statistics and more” and chapter 12 “Predictive analytics with Mahout.”

## **What’s new in the second edition?**

This second edition covers Hadoop 2, which at the time of writing is the current production-ready version of Hadoop. The first edition of the book covered Hadoop 0.22 (Hadoop 1 wasn’t yet out), and Hadoop 2 has turned the world upside-down and opened up the Hadoop platform to processing paradigms beyond MapReduce. YARN, the new scheduler and application manager in Hadoop 2, is complex and new to the community, which prompted me to dedicate a new chapter 2 to covering YARN basics and to discussing how MapReduce now functions as a YARN application.

Parquet has also recently emerged as a new way to store data in HDFS—its columnar format can yield both space and time efficiencies in your data pipelines, and it’s quickly becoming the ubiquitous way to store data. Chapter 4 includes extensive coverage of Parquet, which includes how Parquet supports sophisticated object models such as Avro and how various Hadoop tools can use Parquet.

How data is being ingested into Hadoop has also evolved since the first edition, and Kafka has emerged as the new data pipeline, which serves as the transport tier between your data producers and data consumers, where a consumer would be a system such as Camus that can pull data from Kafka into HDFS. Chapter 5, which covers moving data into and out of Hadoop, now includes coverage of Kafka and Camus.

There are many new technologies that YARN now can support side by side in the same cluster, and some of the more exciting and promising technologies are covered in the new part 4, titled “Beyond MapReduce,” where I cover some compelling new SQL technologies such as Impala and Spark SQL. The last chapter, also new for this edition, looks at how you can write your own YARN application, and it’s packed with information about important features to support your YARN application.

### **Getting help**

You’ll no doubt have many questions when working with Hadoop. Luckily, between the wikis and a vibrant user community, your needs should be well covered:

- The main wiki is located at <http://wiki.apache.org/hadoop/>, and it contains useful presentations, setup instructions, and troubleshooting instructions.
- The Hadoop Common, HDFS, and MapReduce mailing lists can all be found at [http://hadoop.apache.org/mailling\\_lists.html](http://hadoop.apache.org/mailling_lists.html).
- “Search Hadoop” is a useful website that indexes all of Hadoop and its ecosystem projects, and it provides full-text search capabilities: <http://search-hadoop.com/>.
- You’ll find many useful blogs you should subscribe to in order to keep on top of current events in Hadoop. This preface includes a selection of my favorites:
  - Cloudera and Hortonworks are both prolific writers of practical applications on Hadoop—reading their blogs is always educational: <http://www.cloudera.com/blog/> and <http://hortonworks.com/blog/>.
  - Michael Noll is one of the first bloggers to provide detailed setup instructions for Hadoop, and he continues to write about real-life challenges: [www.michael-noll.com/](http://www.michael-noll.com/).
  - There’s a plethora of active Hadoop Twitter users that you may want to follow, including Arun Murthy (@acmurthy), Tom White (@tom\_e\_white), Eric Sammer (@esammer), Doug Cutting (@cutting), and Todd Lipcon (@tlipcon). The Hadoop project tweets on @hadoop.

### **Code conventions and downloads**

All source code in listings or in text is presented in a fixed-width font like this to separate it from ordinary text. Code annotations accompany many of the listings, highlighting important concepts.

All of the text and examples in this book work with Hadoop 2.x, and most of the MapReduce code is written using the newer `org.apache.hadoop.mapreduce` MapReduce APIs. The few examples that use the older `org.apache.hadoop.mapred` package are usually the result of working with a third-party library or a utility that only works with the old API.

All of the code used in this book is available on GitHub at <https://github.com/alexholmes/hiped2> and also from the publisher's website at [www.manning.com/HadoopinPracticeSecondEdition](http://www.manning.com/HadoopinPracticeSecondEdition). The first section in the appendix shows you how to download, install, and get up and running with the code.

### **Third-party libraries**

I use a number of third-party libraries for convenience purposes. They're included in the Maven-built JAR, so there's no extra work required to work with these libraries.

### **Datasets**

Throughout this book, you'll work with three datasets to provide some variety in the examples. All the datasets are small to make them easy to work with. Copies of the exact data used are available in the GitHub repository in the <https://github.com/alexholmes/hiped2/tree/master/test-data> directory. I also sometimes use data that's specific to a chapter, and it's available within chapter-specific subdirectories under the same GitHub location.

### **NASDAQ financial stocks**

I downloaded the NASDAQ daily exchange data from InfoChimps ([www.infochimps.com](http://www.infochimps.com)). I filtered this huge dataset down to just five stocks and their start-of-year values from 2000 through 2009. The data used for this book is available on GitHub at <https://github.com/alexholmes/hiped2/blob/master/test-data/stocks.txt>.

The data is in CSV form, and the fields are in the following order:

```
Symbol,Date,Open,High,Low,Close,Volume,Adj Close
```

### **Apache log data**

I created a sample log file in Apache Common Log Format<sup>1</sup> with some fake Class E IP addresses and some dummy resources and response codes. The file is available on GitHub at <https://github.com/alexholmes/hiped2/blob/master/test-data/apache.log.txt>.

---

<sup>1</sup> See <http://httpd.apache.org/docs/1.3/logs.html#common>.

### **Names**

Names were retrieved from the U.S. government census at [www.census.gov/genealogy/www/data/1990surnames/dist.all.last](http://www.census.gov/genealogy/www/data/1990surnames/dist.all.last), and this data is available at <https://github.com/alexholmes/hiped2/blob/master/test-data/names.txt>.

### **Author Online**

Purchase of *Hadoop in Practice, Second Edition* includes free access to a private web forum run by Manning Publications where you can make comments about the book, ask technical questions, and receive help from the authors and from other users. To access the forum and subscribe to it, point your web browser to [www.manning.com/HadoopinPractice](http://www.manning.com/HadoopinPractice), SecondEdition. This page provides information on how to get on the forum once you are registered, what kind of help is available, and the rules of conduct on the forum. It also provides links to the source code for the examples in the book, errata, and other downloads.

Manning's commitment to our readers is to provide a venue where a meaningful dialog between individual readers and between readers and the author can take place. It is not a commitment to any specific amount of participation on the part of the author, whose contribution to the Author Online forum remains voluntary (and unpaid). We suggest you try asking the author challenging questions lest his interest strays!

The Author Online forum and the archives of previous discussions will be accessible from the publisher's website as long as the book is in print.

## *about the cover illustration*

---

The figure on the cover of *Hadoop in Practice, Second Edition* is captioned “Momak from Kistanja, Dalmatia.” The illustration is taken from a reproduction of an album of traditional Croatian costumes from the mid-nineteenth century by Nikola Arsenovic, published by the Ethnographic Museum in Split, Croatia, in 2003. The illustrations were obtained from a helpful librarian at the Ethnographic Museum in Split, itself situated in the Roman core of the medieval center of the town: the ruins of Emperor Diocletian’s retirement palace from around AD 304. The book includes finely colored illustrations of figures from different regions of Croatia, accompanied by descriptions of the costumes and of everyday life.

Kistanja is a small town located in Bukovica, a geographical region in Croatia. It is situated in northern Dalmatia, an area rich in Roman and Venetian history. The word “momak” in Croatian means a bachelor, beau, or suitor—a single young man who is of courting age—and the young man on the cover, looking dapper in a crisp, white linen shirt and a colorful, embroidered vest, is clearly dressed in his finest clothes, which would be worn to church and for festive occasions—or to go calling on a young lady.

Dress codes and lifestyles have changed over the last 200 years, and the diversity by region, so rich at the time, has faded away. It is now hard to tell apart the inhabitants of different continents, let alone of different hamlets or towns separated by only a few miles. Perhaps we have traded cultural diversity for a more varied personal life—certainly for a more varied and fast-paced technological life.



Manning celebrates the inventiveness and initiative of the computer business with book covers based on the rich diversity of regional life of two centuries ago, brought back to life by illustrations from old books and collections like this one.

# *Part 1*

## *Background and fundamentals*

**P**art 1 of this book consists of chapters 1 and 2, which cover the important Hadoop fundamentals.

Chapter 1 covers Hadoop's components and its ecosystem and provides instructions for installing a pseudo-distributed Hadoop setup on a single host, along with a system that will enable you to run all of the examples in the book. Chapter 1 also covers the basics of Hadoop configuration, and walks you through how to write and run a MapReduce job on your new setup.

Chapter 2 introduces YARN, which is a new and exciting development in Hadoop version 2, transitioning Hadoop from being a MapReduce-only system to one that can support many execution engines. Given that YARN is new to the community, the goal of this chapter is to look at some basics such as its components, how configuration works, and also how MapReduce works as a YARN application. Chapter 2 also provides an overview of some applications that YARN has enabled to execute on Hadoop, such as Spark and Storm.



# *Hadoop in a heartbeat*

---



## ***This chapter covers***

- Examining how the core Hadoop system works
- Understanding the Hadoop ecosystem
- Running a MapReduce job

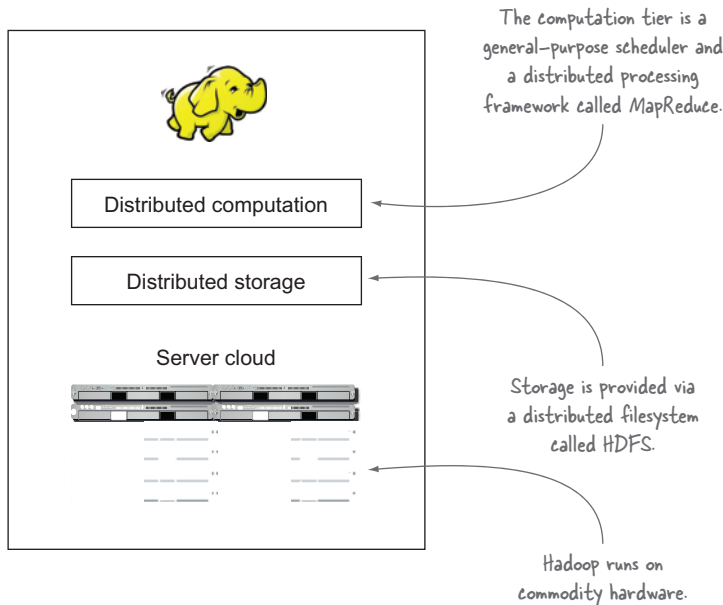
We live in the age of big data, where the data volumes we need to work with on a day-to-day basis have outgrown the storage and processing capabilities of a single host. Big data brings with it two fundamental challenges: how to store and work with voluminous data sizes, and more important, how to understand data and turn it into a competitive advantage.

Hadoop fills a gap in the market by effectively storing and providing computational capabilities for substantial amounts of data. It's a distributed system made up of a distributed filesystem, and it offers a way to parallelize and execute programs on a cluster of machines (see figure 1.1). You've most likely come across Hadoop because it's been adopted by technology giants like Yahoo!, Facebook, and Twitter to address their big data needs, and it's making inroads across all industrial sectors.

Because you've come to this book to get some practical experience with Hadoop and Java,<sup>1</sup> I'll start with a brief overview and then show you how to install

---

<sup>1</sup> To benefit from this book, you should have some practical experience with Hadoop and understand the basic concepts of MapReduce and HDFS (covered in Manning's *Hadoop in Action* by Chuck Lam, 2010). Further, you should have an intermediate-level knowledge of Java—*Effective Java*, 2nd Edition by Joshua Bloch (Addison-Wesley, 2008) is an excellent resource on this topic.



**Figure 1.1** The Hadoop environment is a distributed system that runs on commodity hardware.

Hadoop and run a MapReduce job. By the end of this chapter, you'll have had a basic refresher on the nuts and bolts of Hadoop, which will allow you to move on to the more challenging aspects of working with it.

Let's get started with a detailed overview.

## 1.1 What is Hadoop?

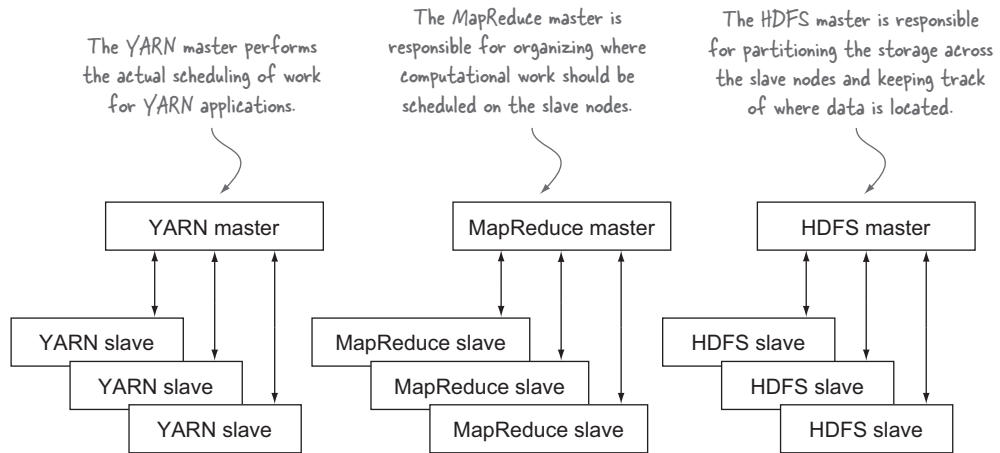
Hadoop is a platform that provides both distributed storage and computational capabilities. Hadoop was first conceived to fix a scalability issue that existed in Nutch,<sup>2</sup> an open source crawler and search engine. At the time, Google had published papers that described its novel distributed filesystem, the Google File System (GFS), and MapReduce, a computational framework for parallel processing. The successful implementation of these papers' concepts in Nutch resulted in it being split into two separate projects, the second of which became Hadoop, a first-class Apache project.

In this section we'll look at Hadoop from an architectural perspective, examine how industry uses it, and consider some of its weaknesses. Once we've covered this background, we'll look at how to install Hadoop and run a MapReduce job.

Hadoop proper, as shown in figure 1.2, is a distributed master-slave architecture<sup>3</sup> that consists of the following primary components:

<sup>2</sup> The Nutch project, and by extension Hadoop, was led by Doug Cutting and Mike Cafarella.

<sup>3</sup> A model of communication where one process, called the *master*, has control over one or more other processes, called *slaves*.



**Figure 1.2** High-level Hadoop 2 master-slave architecture

- Hadoop Distributed File System (HDFS) for data storage.
- Yet Another Resource Negotiator (YARN), introduced in Hadoop 2, a general-purpose scheduler and resource manager. Any YARN application can run on a Hadoop cluster.
- MapReduce, a batch-based computational engine. In Hadoop 2, MapReduce is implemented as a YARN application.

Traits intrinsic to Hadoop are data partitioning and parallel computation of large datasets. Its storage and computational capabilities scale with the addition of hosts to a Hadoop cluster; clusters with hundreds of hosts can easily reach data volumes in the petabytes.

In the first step in this section, we'll examine the HDFS, YARN, and MapReduce architectures.

### 1.1.1 Core Hadoop components

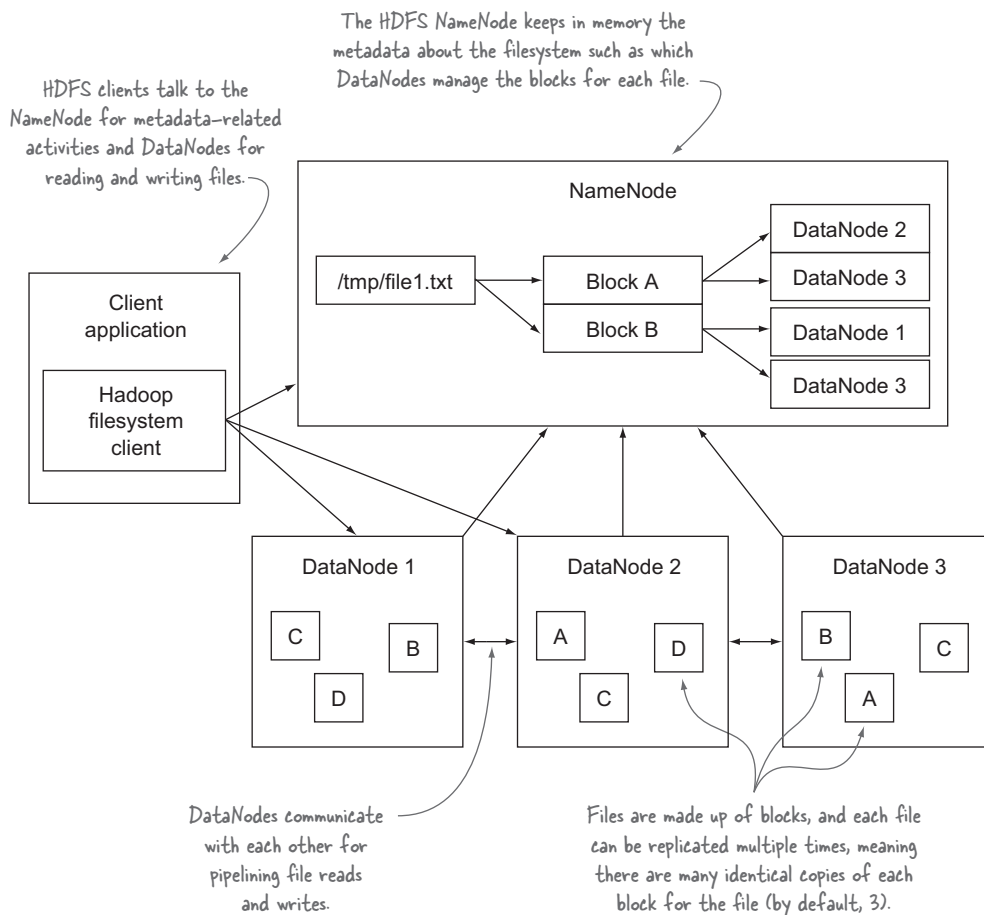
To understand Hadoop's architecture we'll start by looking at the basics of HDFS.

#### HDFS

HDFS is the storage component of Hadoop. It's a distributed filesystem that's modeled after the Google File System (GFS) paper.<sup>4</sup> HDFS is optimized for high throughput and works best when reading and writing large files (gigabytes and larger). To support this throughput, HDFS uses unusually large (for a filesystem) block sizes and data locality optimizations to reduce network input/output (I/O).

Scalability and availability are also key traits of HDFS, achieved in part due to data replication and fault tolerance. HDFS replicates files for a configured number of times, is tolerant of both software and hardware failure, and automatically re-replicates data blocks on nodes that have failed.

<sup>4</sup> See "The Google File System," <http://research.google.com/archive/gfs.html>.



**Figure 1.3** An HDFS client communicating with the master NameNode and slave DataNodes

Figure 1.3 shows a logical representation of the components in HDFS: the NameNode and the DataNode. It also shows an application that’s using the Hadoop filesystem library to access HDFS.

Hadoop 2 introduced two significant new features for HDFS—Federation and High Availability (HA):

- Federation allows HDFS metadata to be shared across multiple NameNode hosts, which aides with HDFS scalability and also provides data isolation, allowing different applications or teams to run their own NameNodes without fear of impacting other NameNodes on the same cluster.
- High Availability in HDFS removes the single point of failure that existed in Hadoop 1, wherein a NameNode disaster would result in a cluster outage. HDFS HA also offers the ability for failover (the process by which a standby NameNode takes over work from a failed primary NameNode) to be automated.

Now that you have a bit of HDFS knowledge, it's time to look at YARN, Hadoop's scheduler.

## YARN

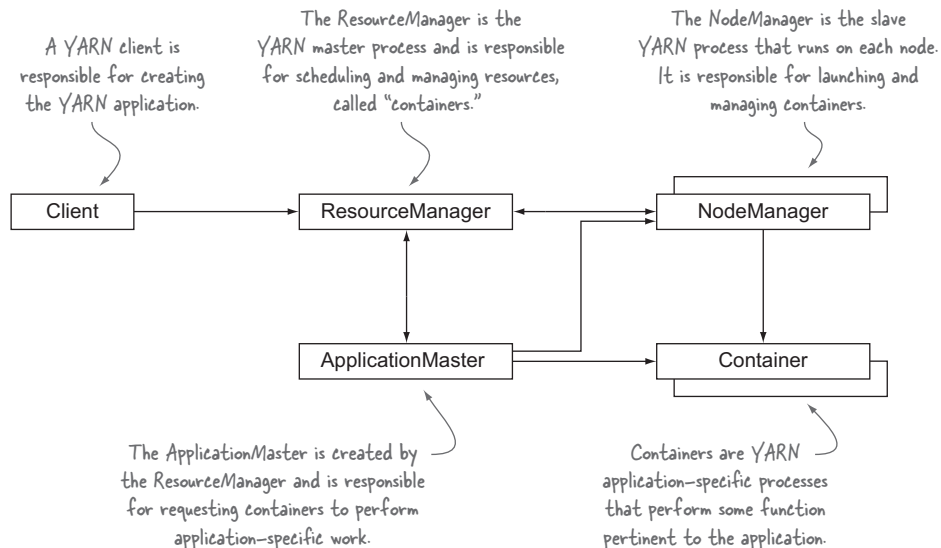
YARN is Hadoop's distributed resource scheduler. YARN is new to Hadoop version 2 and was created to address challenges with the Hadoop 1 architecture:

- Deployments larger than 4,000 nodes encountered scalability issues, and adding additional nodes didn't yield the expected linear scalability improvements.
- Only MapReduce workloads were supported, which meant it wasn't suited to run execution models such as machine learning algorithms that often require iterative computations.

For Hadoop 2 these problems were solved by extracting the scheduling function from MapReduce and reworking it into a generic application scheduler, called YARN. With this change, Hadoop clusters are no longer limited to running MapReduce workloads; YARN enables a new set of workloads to be natively supported on Hadoop, and it allows alternative processing models, such as graph processing and stream processing, to coexist with MapReduce. Chapters 2 and 10 cover YARN and how to write YARN applications.

YARN's architecture is simple because its primary role is to schedule and manage resources in a Hadoop cluster. Figure 1.4 shows a logical representation of the core components in YARN: the ResourceManager and the NodeManager. Also shown are the components specific to YARN applications, namely, the YARN application client, the ApplicationMaster, and the container.

To fully realize the dream of a generalized distributed platform, Hadoop 2 introduced another change—the ability to allocate containers in various configurations.



**Figure 1.4** The logical YARN architecture showing typical communication between the core YARN components and YARN application components



Hadoop 1 had the notion of “slots,” which were a fixed number of map and reduce processes that were allowed to run on a single node. This was wasteful in terms of cluster utilization and resulted in underutilized resources during MapReduce operations, and it also imposed memory limits for map and reduce tasks. With YARN, each container requested by an ApplicationMaster can have disparate memory and CPU traits, and this gives YARN applications full control over the resources they need to fulfill their work.

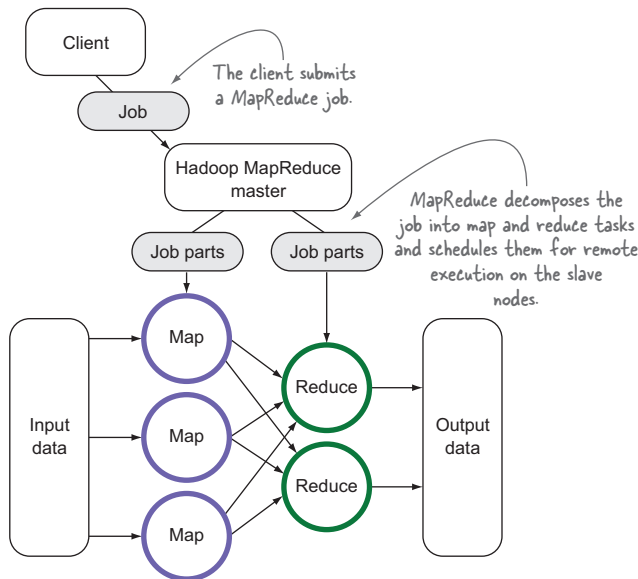
You’ll work with YARN in more detail in chapters 2 and 10, where you’ll learn how YARN works and how to write a YARN application. Next up is an examination of MapReduce, Hadoop’s computation engine.

### MAPREDUCE

MapReduce is a batch-based, distributed computing framework modeled after Google’s paper on MapReduce.<sup>5</sup> It allows you to parallelize work over a large amount of raw data, such as combining web logs with relational data from an OLTP database to model how users interact with your website. This type of work, which could take days or longer using conventional serial programming techniques, can be reduced to minutes using MapReduce on a Hadoop cluster.

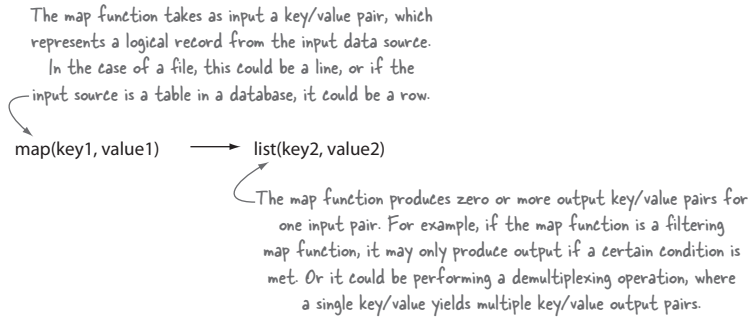
The MapReduce model simplifies parallel processing by abstracting away the complexities involved in working with distributed systems, such as computational parallelization, work distribution, and dealing with unreliable hardware and software. With this abstraction, MapReduce allows the programmer to focus on addressing business needs rather than getting tangled up in distributed system complications.

MapReduce decomposes work submitted by a client into small parallelized map and reduce tasks, as shown in figure 1.5. The map and reduce constructs used in



**Figure 1.5** A client submitting a job to MapReduce, breaking the work into small map and reduce tasks

<sup>5</sup> See “MapReduce: Simplified Data Processing on Large Clusters,” <http://research.google.com/archive/mapreduce.html>.

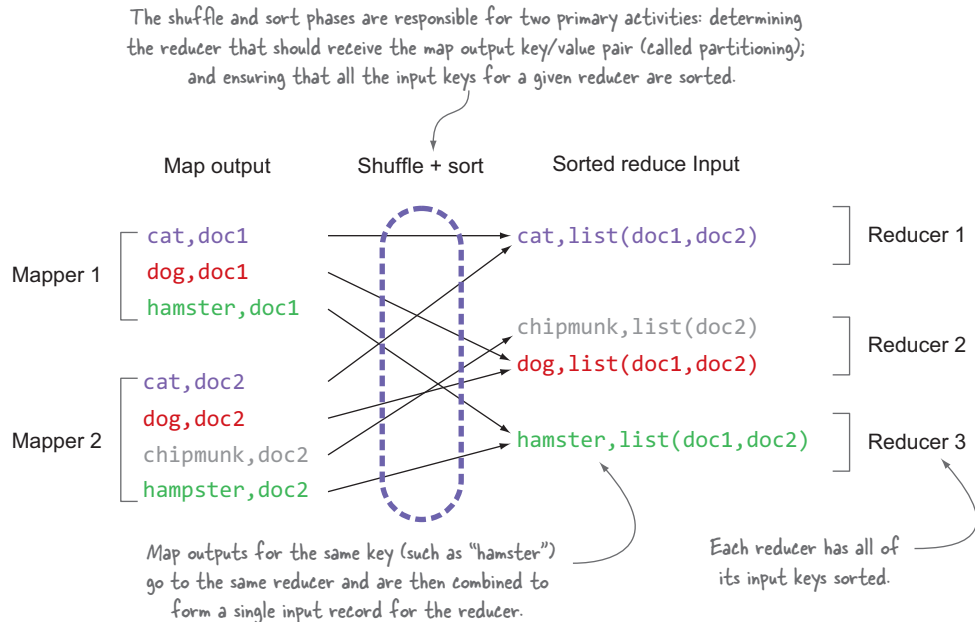


**Figure 1.6** A logical view of the map function that takes a key/value pair as input

MapReduce are borrowed from those found in the Lisp functional programming language, and they use a shared-nothing model to remove any parallel execution interdependencies that could add unwanted synchronization points or state sharing.<sup>6</sup>

The role of the programmer is to define map and reduce functions where the map function outputs key/value tuples, which are processed by reduce functions to produce the final output. Figure 1.6 shows a pseudocode definition of a map function with regard to its input and output.

The power of MapReduce occurs between the map output and the reduce input in the shuffle and sort phases, as shown in figure 1.7.



**Figure 1.7** MapReduce's shuffle and sort phases

<sup>6</sup> A shared-nothing architecture is a distributed computing concept that represents the notion that each node is independent and self-sufficient.