

MongoDB

IN ACTION

SECOND EDITION

Kyle Banker
Peter Bakkum
Shaun Verch
Douglas Garrett
Tim Hawkins

 MANNING



Covers MongoDB version 3.0

www.it-ebooks.info

MongoDB in Action

www.it-ebooks.info

www.it-ebooks.info

MongoDB in Action

Second Edition

KYLE BANKER
PETER BAKKUM
SHAUN VERCH
DOUGLAS GARRETT
TIM HAWKINS



MANNING
SHELTER ISLAND

www.it-ebooks.info

For online information and ordering of this and other Manning books, please visit 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

©2016 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
PO Box 761
Shelter Island, NY 11964

Development editors: Susan Conant, Jeff Bleiel
Technical development editors: Brian Hanafée, Jürgen Hoffman,
Wouter Thielen
Copyeditors: Liz Welch, Jodie Allen
Proofreader: Melody Dolab
Technical proofreader: Doug Warren
Typesetter: Dennis Dalinnik
Cover designer: Marija Tudor

ISBN: 9781617291609

Printed in the United States of America

1 2 3 4 5 6 7 8 9 10 – EBM – 21 20 19 18 17 16

www.it-ebooks.info

*This book is dedicated to peace and human dignity
and to all those who work for these ideals*

www.it-ebooks.info

brief contents

PART 1	GETTING STARTED	1
1	■ A database for the modern web	3
2	■ MongoDB through the JavaScript shell	29
3	■ Writing programs using MongoDB	52
PART 2	APPLICATION DEVELOPMENT IN MONGODB.....	71
4	■ Document-oriented data	73
5	■ Constructing queries	98
6	■ Aggregation	120
7	■ Updates, atomic operations, and deletes	157
PART 3	MONGODB MASTERY.....	195
8	■ Indexing and query optimization	197
9	■ Text search	244
10	■ WiredTiger and pluggable storage	273
11	■ Replication	296
12	■ Scaling your system with sharding	333
13	■ Deployment and administration	376

contents

preface xvii
acknowledgments xix
about this book xxi
about the cover illustration xxiv

PART 1 GETTING STARTED.1

1 A database for the modern web 3

- 1.1 Built for the internet 5
- 1.2 MongoDB's key features 6
 - Document data model 6* ▪ *Ad hoc queries 10*
 - Indexes 10* ▪ *Replication 11* ▪ *Speed and durability 12*
 - Scaling 14*
- 1.3 MongoDB's core server and tools 15
 - Core server 16* ▪ *JavaScript shell 16* ▪ *Database drivers 17*
 - Command-line tools 18*
- 1.4 Why MongoDB? 18
 - MongoDB versus other databases 19* ▪ *Use cases and production deployments 22*
- 1.5 Tips and limitations 24
- 1.6 History of MongoDB 25

- 1.7 Additional resources 27
- 1.8 Summary 28

2 *MongoDB through the JavaScript shell* 29

- 2.1 Diving into the MongoDB shell 30
 - Starting the shell* 30 ▪ *Databases, collections, and documents* 31
 - Inserts and queries* 32 ▪ *Updating documents* 34
 - Deleting data* 38 ▪ *Other shell features* 38
- 2.2 Creating and querying with indexes 39
 - Creating a large collection* 39 ▪ *Indexing and explain()* 41
- 2.3 Basic administration 46
 - Getting database information* 46 ▪ *How commands work* 48
- 2.4 Getting help 49
- 2.5 Summary 51

3 *Writing programs using MongoDB* 52

- 3.1 MongoDB through the Ruby lens 53
 - Installing and connecting* 53 ▪ *Inserting documents in Ruby* 55
 - Queries and cursors* 56 ▪ *Updates and deletes* 57
 - Database commands* 58
- 3.2 How the drivers work 59
 - Object ID generation* 59
- 3.3 Building a simple application 61
 - Setting up* 61 ▪ *Gathering data* 62 ▪ *Viewing the archive* 65
- 3.4 Summary 69

PART 2 APPLICATION DEVELOPMENT IN MONGODB71

4 *Document-oriented data* 73

- 4.1 Principles of schema design 74
- 4.2 Designing an e-commerce data model 75
 - Schema basics* 76 ▪ *Users and orders* 80 ▪ *Reviews* 83
- 4.3 Nuts and bolts: On databases, collections, and documents 84
 - Databases* 84 ▪ *Collections* 87 ▪ *Documents and insertion* 92
- 4.4 Summary 96

-
- ## 5 *Constructing queries* 98
- 5.1 E-commerce queries 99
 - Products, categories, and reviews* 99 ▪ *Users and orders* 101
 - 5.2 MongoDB's query language 103
 - Query criteria and selectors* 103 ▪ *Query options* 117
 - 5.3 Summary 119
- ## 6 *Aggregation* 120
- 6.1 Aggregation framework overview 121
 - 6.2 E-commerce aggregation example 123
 - Products, categories, and reviews* 125
 - User and order* 132
 - 6.3 Aggregation pipeline operators 135
 - \$project* 136 ▪ *\$group* 136 ▪ *\$match, \$sort, \$skip, \$limit* 138 ▪ *\$unwind* 139 ▪ *\$out* 139
 - 6.4 Reshaping documents 140
 - String functions* 141 ▪ *Arithmetic functions* 142
 - Date functions* 142 ▪ *Logical functions* 143
 - Set Operators* 144 ▪ *Miscellaneous functions* 145
 - 6.5 Understanding aggregation pipeline performance 146
 - Aggregation pipeline options* 147 ▪ *The aggregation framework's explain() function* 147 ▪ *allowDiskUse option* 151
 - Aggregation cursor option* 151
 - 6.6 Other aggregation capabilities 152
 - .count() and .distinct()* 153 ▪ *map-reduce* 153
 - 6.7 Summary 156
- ## 7 *Updates, atomic operations, and deletes* 157
- 7.1 A brief tour of document updates 158
 - Modify by replacement* 159 ▪ *Modify by operator* 159
 - Both methods compared* 160 ▪ *Deciding: replacement vs. operators* 160
 - 7.2 E-commerce updates 162
 - Products and categories* 162 ▪ *Reviews* 167 ▪ *Orders* 168
 - 7.3 Atomic document processing 171
 - Order state transitions* 172 ▪ *Inventory management* 174

- 7.4 Nuts and bolts: MongoDB updates and deletes 179
 - Update types and options* 179
 - *Update operators* 181
 - The findAndModify command* 188
 - *Deletes* 189
 - Concurrency, atomicity, and isolation* 190
 - Update performance notes* 191
- 7.5 Reviewing update operators 192
- 7.6 Summary 193

PART 3 MONGODB MASTERY195

8 *Indexing and query optimization* 197

- 8.1 Indexing theory 198
 - A thought experiment* 198
 - *Core indexing concepts* 201
 - B-trees* 205
- 8.2 Indexing in practice 207
 - Index types* 207
 - *Index administration* 211
- 8.3 Query optimization 216
 - Identifying slow queries* 217
 - *Examining slow queries* 221
 - Query patterns* 241
- 8.4 Summary 243

9 *Text search* 244

- 9.1 Text searches—not just pattern matching 245
 - Text searches vs. pattern matching* 246
 - *Text searches vs. web page searches* 247
 - *MongoDB text search vs. dedicated text search engines* 250
- 9.2 Manning book catalog data download 253
- 9.3 Defining text search indexes 255
 - Text index size* 255
 - *Assigning an index name and indexing all text fields in a collection* 256
- 9.4 Basic text search 257
 - More complex searches* 259
 - *Text search scores* 261
 - Sorting results by text search score* 262
- 9.5 Aggregation framework text search 263
 - Where's MongoDB in Action, Second Edition?* 265

- 9.6 Text search languages 267
 - Specifying language in the index* 267
 - *Specifying the language in the document* 269
 - *Specifying the language in a search* 269
 - Available languages* 271
- 9.7 Summary 272

10 **WiredTiger and pluggable storage** 273

- 10.1 Pluggable Storage Engine API 273
 - Why use different storages engines?* 274
- 10.2 WiredTiger 275
 - Switching to WiredTiger* 276
 - *Migrating your database to WiredTiger* 277
- 10.3 Comparison with MMAPv1 278
 - Configuration files* 279
 - *Insertion script and benchmark script* 281
 - *Insertion benchmark results* 283
 - Read performance scripts* 285
 - *Read performance results* 286
 - Benchmark conclusion* 288
- 10.4 Other examples of pluggable storage engines 289
- 10.5 Advanced topics 290
 - How does a pluggable storage engine work?* 290
 - Data structure* 292
 - *Locking* 294
- 10.6 Summary 295

11 **Replication** 296

- 11.1 Replication overview 297
 - Why replication matters* 297
 - *Replication use cases and limitations* 298
- 11.2 Replica sets 300
 - Setup* 300
 - *How replication works* 307
 - Administration* 314
- 11.3 Drivers and replication 324
 - Connections and failover* 324
 - *Write concern* 327
 - Read scaling* 328
 - *Tagging* 330
- 11.4 Summary 332

12 *Scaling your system with sharding* 333

- 12.1 Sharding overview 334
 - What is sharding?* 334
 - When should you shard?* 335
- 12.2 Understanding components of a sharded cluster 336
 - Shards: storage of application data* 337
 - Mongos router: router of operations* 338
 - Config servers: storage of metadata* 338
- 12.3 Distributing data in a sharded cluster 339
 - Ways data can be distributed in a sharded cluster* 340
 - Distributing databases to shards* 341
 - Sharding within collections* 341
- 12.4 Building a sample shard cluster 343
 - Starting the mongod and mongos servers* 343
 - Configuring the cluster* 346
 - Sharding collections* 347
 - Writing to a sharded cluster* 349
- 12.5 Querying and indexing a shard cluster 355
 - Query routing* 355
 - Indexing in a sharded cluster* 356
 - The explain() tool in a sharded cluster* 357
 - Aggregation in a sharded cluster* 359
- 12.6 Choosing a shard key 359
 - Imbalanced writes (hotspots)* 360
 - Unsplittable chunks (coarse granularity)* 362
 - Poor targeting (shard key not present in queries)* 362
 - Ideal shard keys* 363
 - Inherent design trade-offs (email application)* 364
- 12.7 Sharding in production 365
 - Provisioning* 366
 - Deployment* 369
 - Maintenance* 370
- 12.8 Summary 375

13 *Deployment and administration* 376

- 13.1 Hardware and provisioning 377
 - Cluster topology* 377
 - Deployment environment* 378
 - Provisioning* 385
- 13.2 Monitoring and diagnostics 386
 - Logging* 387
 - MongoDB diagnostic commands* 387
 - MongoDB diagnostic tools* 388
 - MongoDB Monitoring Service* 390
 - External monitoring applications* 390
- 13.3 Backups 391
 - mongodump and mongorestore* 391
 - Data file-based backups* 392
 - MMS backups* 393

13.4	Security	394
	<i>Secure environments</i>	394
	<i>Network encryption</i>	395
	<i>Authentication</i>	397
	▪ <i>Replica set authentication</i>	401
	▪ <i>Sharding authentication</i>	402
	▪ <i>Enterprise security features</i>	402
13.5	Administrative tasks	402
	<i>Data imports and exports</i>	402
	▪ <i>Compaction and repair</i>	403
	▪ <i>Upgrading</i>	405
13.6	Performance troubleshooting	405
	▪ <i>Working set</i>	406
	▪ <i>Performance cliff</i>	407
	▪ <i>Query interactions</i>	407
	▪ <i>Seek professional assistance</i>	408
13.7	Deployment checklist	408
13.8	Summary	410
<i>appendix A</i>	<i>Installation</i>	<i>411</i>
<i>appendix B</i>	<i>Design patterns</i>	<i>421</i>
<i>appendix C</i>	<i>Binary data and GridFS</i>	<i>433</i>
	<i>index</i>	<i>441</i>

preface

Databases are the workhorses of the information age. Like Atlas, they go largely unnoticed in supporting the digital world we've come to inhabit. It's easy to forget that our digital interactions, from commenting and tweeting to searching and sorting, are in essence interactions with a database. Because of this fundamental yet hidden function, I always experience a certain sense of awe when thinking about databases, not unlike the awe one might feel when walking across a suspension bridge normally reserved for automobiles.

The database has taken many forms. The indexes of books and the card catalogs that once stood in libraries are both databases of a sort, as are the ad hoc structured text files of the Perl programmers of yore. Perhaps most recognizable now as databases proper are the sophisticated, fortune-making relational databases that underlie much of the world's software. These relational databases, with their idealized third-normal forms and expressive SQL interfaces, still command the respect of the old guard, and appropriately so.

But as a working web application developer a few years back, I was eager to sample the emerging alternatives to the reigning relational database. When I discovered MongoDB, the resonance was immediate. I liked the idea of using a JSON-like structure to represent data. JSON is simple, intuitive, and human-friendly. That MongoDB also based its query language on JSON lent a high degree of comfort and harmony to the usage of this new database. The interface came first. Compelling features like easy replication and sharding made the package all the more intriguing. And by the time

I'd built a few applications on MongoDB and beheld the ease of development it imparted, I'd become a convert.

Through an unlikely turn of events, I started working for 10gen, the company spearheading the development of this open source database. For two years, I've had the opportunity to improve various client drivers and work with numerous customers on their MongoDB deployments. The experience gained through this process has, I hope, been distilled faithfully into the book you're reading now.

As a piece of software and a work in progress, MongoDB is still far from perfection. But it's also successfully supporting thousands of applications atop database clusters small and large, and it's maturing daily. It's been known to bring out wonder, even happiness, in many a developer. My hope is that it can do the same for you.

This is the second edition of MongoDB in Action and I hope that you enjoy reading the book!

KYLE BANKER

acknowledgments

Thanks are due to folks at Manning for helping make this book a reality. Michael Stephens helped conceive the first edition of this book, and my development editors for this second edition, Susan Conant, Jeff Bleiel, and Maureen Spencer, pushed the book to completion while being helpful along the way. My thanks go to them.

Book writing is a time-consuming enterprise. I feel I wouldn't have found the time to finish this book had it not been for the generosity of Eliot Horowitz and Dwight Merriman. Eliot and Dwight, through their initiative and ingenuity, created MongoDB, and they trusted me to document the project. My thanks to them.

Many of the ideas in this book owe their origins to conversations I had with colleagues at 10gen. In this regard, special thanks are due to Mike Dirolf, Scott Hernandez, Alvin Richards, and Mathias Stearn. I'm especially indebted to Kristina Chowdorow, Richard Kreuter, and Aaron Staple for providing expert reviews of entire chapters for the first edition.

The following reviewers read the manuscript of the first edition at various stages during its development: Kevin Jackson, Hardy Ferentschik, David Sinclair, Chris Chandler, John Nunemaker, Robert Hanson, Alberto Lerner, Rick Wagner, Ryan Cox, Andy Brudtkuhl, Daniel Bretoi, Greg Donald, Sean Reilly, Curtis Miller, Sanchet Dighe, Philip Hallstrom, and Andy Dingley. And I am also indebted to all the reviewers who read the second edition, including Agustin Treceno, Basheeruddin Ahmed, Gavin Whyte, George Girton, Gregor Zurowski, Hardy Ferentschik, Hernan Garcia, Jeet Marwah, Johan Mattisson, Jonathan Thoms, Julia Varigina, Jürgen Hoffmann, Mike Frey, Phlippie Smith, Scott Lyons, and Steve Johnson. Special thanks go to Wouter Thielen for his work on chapter 10, technical editor Mihalis Tsoukalos, who devoted

many hours to whipping the second edition into shape, and to Doug Warren for his thorough technical review of the second edition shortly before it went to press.

My amazing wife, Dominika, offered her patience and support, through the writing of both editions of this book, and to my wonderful son, Oliver, just for being awesome.

KYLE BANKER

about this book

This book is for application developers and DBAs wanting to learn MongoDB from the ground up. If you're new to MongoDB, you'll find in this book a tutorial that moves at a comfortable pace. If you're already a user, the more detailed reference sections in the book will come in handy and should fill any gaps in your knowledge. In terms of depth, the material should be suitable for all but the most advanced users. Although the book is about the latest MongoDB version, which at the time of writing is 3.0.x, it also covers the previous stable MongoDB version that is 2.6.

The code examples are written in JavaScript, the language of the MongoDB shell, and Ruby, a popular scripting language. Every effort has been made to provide simple but useful examples, and only the plainest features of the JavaScript and Ruby languages are used. The main goal is to present the MongoDB API in the most accessible way possible. If you have experience with other programming languages, you should find the examples easy to follow.

One more note about languages. If you're wondering, "Why couldn't this book use language X?" you can take heart. The officially supported MongoDB drivers feature consistent and analogous APIs. This means that once you learn the basic API for one driver, you can pick up the others fairly easily.

How to use this book

This book is part tutorial, part reference. If you're brand-new to MongoDB, then reading through the book in order makes a lot of sense. There are numerous code examples that you can run on your own to help solidify the concepts. At minimum, you'll

need to install MongoDB and optionally the Ruby driver. Instructions for these installations can be found in appendix A.

If you've already used MongoDB, then you may be more interested in particular topics. Chapters 8 to 13 and all of the appendixes stand on their own and can safely be read in any order. Additionally, chapters 4 to 7 contain the so-called "nuts and bolts" sections, which focus on fundamentals. These also can be read outside the flow of the surrounding text.

Roadmap

This book is divided into three parts.

Part 1 is an end-to-end introduction to MongoDB. Chapter 1 gives an overview of MongoDB's history, features, and use cases. Chapter 2 teaches the database's core concepts through a tutorial on the MongoDB command shell. Chapter 3 walks through the design of a simple application that uses MongoDB on the back end.

Part 2 is an elaboration on the MongoDB API presented in part 1. With a specific focus on application development, the four chapters in part 2 progressively describe a schema and its operations for an e-commerce app. Chapter 4 delves into documents, the smallest unit of data in MongoDB, and puts forth a basic e-commerce schema design. Chapters 5, 6, and 7 then teach you how to work with this schema by covering queries and updates. To augment the presentation, each of the chapters in part 2 contains a detailed breakdown of its subject matter.

Part 3 focuses on MongoDB mastery. Chapter 8 is a thorough study of indexing and query optimization. The subject of Chapter 9 is text searching inside MongoDB. Chapter 10, which is totally new in this edition, is about the WiredTiger storage engine and pluggable storage, which are unique features of MongoDB v3. Chapter 11 concentrates on replication, with strategies for deploying MongoDB for high availability and read scaling. Chapter 12 describes sharding, MongoDB's path to horizontal scalability. And chapter 13 provides a series of best practices for deploying, administering, and troubleshooting MongoDB installations.

The book ends with three appendixes. Appendix A covers installation of MongoDB and Ruby (for the driver examples) on Linux, Mac OS X, and Windows. Appendix B presents a series of schema and application design patterns, and it also includes a list of anti-patterns. Appendix C shows how to work with binary data in MongoDB and how to use GridFS, a spec implemented by all the drivers, to store especially large files in the database.

Code conventions and downloads

All source code in the listings and in the text is presented in a fixed-width font, which separates it from ordinary text.

Code annotations accompany some of the listings, highlighting important concepts. In some cases, numbered bullets link to explanations that follow in the text.

As an open source project, 10gen keeps MongoDB's bug tracker open to the community at large. At several points in the book, particularly in the footnotes, you'll see references to bug reports and planned improvements. For example, the ticket for adding full-text search to the database is SERVER-380. To view the status of any such ticket, point your browser to <http://jira.mongodb.org>, and enter the ticket ID in the search box.

You can download the book's source code, with some sample data, from the book's site at <http://mongodb-book.com> as well as from the publisher's website at <http://manning.com/MongoDBinAction>.

Software requirements

To get the most out of this book, you'll need to have MongoDB installed on your system. Instructions for installing MongoDB can be found in appendix A and also on the official MongoDB website (<http://mongodb.org>).

If you want to run the Ruby driver examples, you'll also need to install Ruby. Again, consult appendix A for instructions on this.

Author Online

The purchase of *MongoDB in Action, Second Edition* includes free access to a private forum run by Manning Publications where you can make comments about the book, ask technical questions, and receive help from the author and other users. To access and subscribe to the forum, point your browser to www.manning.com/MongoDBinAction. 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 in the forum.

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

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.

- [*click Blog Theory: Feedback and Capture in the Circuits of Drive*](#)
- [Gilchrist on Blake: The Life of William Blake by Alexander Gilchrist \(Lives That Never Grow Old\) pdf, azw \(kindle\)](#)
- [I Sank the Bismarck: Memoirs of a Second World War Navy Pilot here](#)
- [click Photography and Its Violations \(Columbia Themes in Philosophy, Social Criticism, and the Arts\)](#)
- [download The Real History Behind the Da Vinci Code pdf, azw \(kindle\), epub](#)
- [read online Bread: The very best recipes for loaves, rolls, knots and twists from around the world](#)

- <http://www.netc-bd.com/ebooks/Red-Rain--A-Novel.pdf>
- <http://betsy.wesleychapelcomputerrepair.com/library/Gilchrist-on-Blake--The-Life-of-William-Blake-by-Alexander-Gilchrist--Lives-That-Never-Grow-Old-.pdf>
- <http://xn--d1aboelcb1f.xn--p1ai/lib/I-Sank-the-Bismarck--Memoirs-of-a-Second-World-War-Navy-Pilot.pdf>
- <http://ramazotti.ru/library/Photography-and-Its-Violations--Columbia-Themes-in-Philosophy--Social-Criticism--and-the-Arts-.pdf>
- <http://monkeybubblemedia.com/lib/Healing-Herbs-A-to-Z--A-Handy-Reference-to-Healing-Plants.pdf>
- <http://www.netc-bd.com/ebooks/Deadly-Curiosities--Deadly-Curiosities--Book-1-.pdf>