

# Table of Contents

19	Limits, Skips, and Sorts	67
20	Immortal Cursors	77
<hr/>		
25	Part II. Designing Your Application	
25	Indexes	77
27	Introduction to Indexes	77
28	Creating an Index	87
29	Introduction to Compound Indexes	18
29	Using Compound Indexes	43
29	How \$ Operators Use Indexes	58
33	Preface	xv
<hr/>		
35	Part I. Introduction to MongoDB	
35	1. Introduction	3
37	Ease of Use	3
40	Designed to Scale	4
49	Rich with Features...	5
49	...Without Sacrificing Speed	6
53	The Philosophy	6
53	2. Getting Started	7
54	Documents	7
55	Collections	8
55	Dynamic Schemas	8
55	Naming	9
56	Databases	10
57	Getting and Starting MongoDB	11
57	Introduction to the MongoDB Shell	13
57	Running the Shell	13
58	A MongoDB Client	14
59	Basic Operations with the Shell	14
63	Data Types	16
65	Basic Data Types	16
66	Dates	18



Arrays	19
Embedded Documents	19
_id and ObjectIds	20
Using the MongoDB Shell	22
Tips for Using the Shell	22
Running Scripts with the Shell	23
Creating a .mongorc.js	25
Customizing Your Prompt	26
Editing Complex Variables	27
Inconvenient Collection Names	28
<b>3. Creating, Updating, and Deleting Documents.....</b>	<b>29</b>
Inserting Documents	29
insertMany	29
Insert Validation	32
insert	33
Removing Documents	33
drop	34
Updating Documents	35
Document Replacement	35
Using Update Operators	37
Upserts	46
Updating Multiple Documents	49
Returning Updated Documents	49
<b>4. Querying.....</b>	<b>53</b>
Introduction to find	53
Specifying Which Keys to Return	54
Limitations	55
Query Criteria	55
Query Conditionals	55
OR Queries	56
\$not	57
Type-Specific Queries	57
null	57
Regular Expressions	58
Querying Arrays	59
Querying on Embedded Documents	63
\$where Queries	65
Cursors	66



Limits, Skips, and Sorts	67
Avoiding Large Skips	68
Immortal Cursors	70

---

## Part II. Designing Your Application

<b>5. Indexes.....</b>	<b>75</b>
Introduction to Indexes	75
Creating an Index	78
Introduction to Compound Indexes	81
How MongoDB Selects an Index	84
Using Compound Indexes	85
How \$ Operators Use Indexes	104
Indexing Objects and Arrays	114
Index Cardinality	116
explain Output	116
When Not to Index	125
Types of Indexes	126
Unique Indexes	126
Partial Indexes	128
Index Administration	129
Identifying Indexes	130
Changing Indexes	130
<b>6. Special Index and Collection Types.....</b>	<b>133</b>
Geospatial Indexes	133
Types of Geospatial Queries	134
Using Geospatial Indexes	136
Compound Geospatial Indexes	144
2d Indexes	144
Indexes for Full Text Search	146
Creating a Text Index	147
Text Search	148
Optimizing Full-Text Search	151
Searching in Other Languages	151
Capped Collections	151
Creating Capped Collections	154
Tailable Cursors	154
Time-To-Live Indexes	155



67	Storing Files with GridFS	156
68	Getting Started with GridFS: mongofiles	156
70	Working with GridFS from the MongoDB Drivers	157
	Under the Hood	158
<hr/>		
<b>7.</b>	<b>Introduction to the Aggregation Framework.....</b>	<b>161</b>
	Pipelines, Stages, and Tunables	161
75	Getting Started with Stages: Familiar Operations .....	163
75	Expressions	168
78	\$project	169
81	\$unwind	174
84	Array Expressions	181
85	Accumulators	186
104	Using Accumulators in Project Stages	186
114	Introduction to Grouping	188
116	The <code>_id</code> Field in Group Stages	192
116	Group Versus Project	195
125	Writing Aggregation Pipeline Results to a Collection	198
<b>8.</b>	<b>Transactions.....</b>	<b>199</b>
128	Introduction to Transactions	199
129	A Definition of ACID	200
130	How to Use Transactions	200
130	Tuning Transaction Limits for Your Application	205
	Timing and Olog Size Limits	205
<b>9.</b>	<b>Application Design.....</b>	<b>207</b>
134	Schema Design Considerations	207
136	Schema Design Patterns	208
144	Normalization Versus Denormalization	211
144	Examples of Data Representations	212
146	Cardinality	216
147	Friends, Followers, and Other Inconveniences	216
148	Optimizations for Data Manipulation	219
151	Removing Old Data	219
151	Planning Out Databases and Collections	220
151	Managing Consistency	221
154	Migrating Schemas	222
154	Managing Schemas	223
155	When Not to Use MongoDB	223



## Part III. Replication

<b>10. Setting Up a Replica Set.....</b>	<b>227</b>
Introduction to Replication	227
Setting Up a Replica Set, Part 1	228
Networking Considerations	229
Security Considerations	230
Setting Up a Replica Set, Part 2	230
Observing Replication	233
Changing Your Replica Set Configuration	238
How to Design a Set	241
How Elections Work	243
Member Configuration Options	244
Priority	244
Hidden Members	245
Election Arbiters	246
Building Indexes	247
<b>11. Components of a Replica Set.....</b>	<b>249</b>
Syncing	249
Initial Sync	251
Replication	253
Handling Staleness	253
Heartbeats	253
Member States	254
Elections	255
Rollbacks	255
When Rollbacks Fail	259
<b>12. Connecting to a Replica Set from Your Application.....</b>	<b>261</b>
Client-to-Replica Set Connection Behavior	261
Waiting for Replication on Writes	263
Other Options for “w”	265
Custom Replication Guarantees	265
Guaranteeing One Server per Data Center	265
Guaranteeing a Majority of Nonhidden Members	267
Creating Other Guarantees	267
Sending Reads to Secondaries	268
Consistency Considerations	268



Load Considerations	269
Reasons to Read from Secondaries	269
<b>13. Administration</b>	<b>271</b>
Starting Members in Standalone Mode	271
Replica Set Configuration	272
Creating a Replica Set	272
Changing Set Members	273
Creating Larger Sets	274
Forcing Reconfiguration	274
Manipulating Member State	275
Turning Primaries into Secondaries	275
Preventing Elections	275
Monitoring Replication	275
Getting the Status	276
Visualizing the Replication Graph	279
Replication Loops	280
Disabling Chaining	281
Calculating Lag	281
Resizing the Oplog	282
Building Indexes	283
Replication on a Budget	285
<hr/>	
<b>Part IV. Sharding</b>	
<b>14. Introduction to Sharding</b>	<b>289</b>
What Is Sharding?	289
Understanding the Components of a Cluster	290
Sharding on a Single-Machine Cluster	291
<b>15. Configuring Sharding</b>	<b>303</b>
When to Shard	303
Starting the Servers	304
Config Servers	304
The mongos Processes	305
Adding a Shard from a Replica Set	306
Adding Capacity	310
Sharding Data	310
How MongoDB Tracks Cluster Data	311



Chunk Ranges	312
Splitting Chunks	314
The Balancer	316
Collations	317
Change Streams	317
<b>16. Choosing a Shard Key.....</b>	<b>319</b>
Taking Stock of Your Usage	319
Picturing Distributions	320
Ascending Shard Keys	320
Randomly Distributed Shard Keys	323
Location-Based Shard Keys	325
Shard Key Strategies	327
Hashed Shard Key	327
Hashed Shard Keys for GridFS	328
The Firehose Strategy	329
Multi-Hotspot	330
Shard Key Rules and Guidelines	334
Shard Key Limitations	334
Shard Key Cardinality	334
Controlling Data Distribution	334
Using a Cluster for Multiple Databases and Collections	335
Manual Sharding	336
<b>17. Sharding Administration.....</b>	<b>339</b>
Seeing the Current State	339
Getting a Summary with sh.status()	339
Seeing Configuration Information	341
Tracking Network Connections	348
Getting Connection Statistics	348
Limiting the Number of Connections	354
Server Administration	356
Adding Servers	356
Changing Servers in a Shard	356
Removing a Shard	356
Balancing Data	359
The Balancer	360
Changing Chunk Size	361
Moving Chunks	362
Jumbo Chunks	364



312	Refreshing Configurations	367
<hr/>		
<b>Part V. Application Administration</b>		
<b>18.</b>	<b>Seeing What Your Application Is Doing</b>	<b>371</b>
	Seeing the Current Operations	371
	Finding Problematic Operations	374
	Killing Operations	375
	False Positives	375
	Preventing Phantom Operations	375
	Using the System Profiler	376
	Calculating Sizes	379
	Documents	379
	Collections	380
	Databases	385
	Using mongotop and mongostat	386
<b>19.</b>	<b>An Introduction to MongoDB Security</b>	<b>389</b>
	MongoDB Authentication and Authorization	389
	Authentication Mechanisms	389
	Authorization	390
	Using x.509 Certificates to Authenticate Both Members and Clients	392
	A Tutorial on MongoDB Authentication and Transport Layer Encryption	395
	Establish a CA	395
	Generate and Sign Member Certificates	400
	Generate and Sign Client Certificates	401
	Bring Up the Replica Set Without Authentication and Authorization Enabled	401
	Create the Admin User	402
	Restart the Replica Set with Authentication and Authorization Enabled	403
<b>20.</b>	<b>Durability</b>	<b>405</b>
	Durability at the Member Level Through Journaling	405
	Durability at the Cluster Level Using Write Concern	407
	The w and wtimeout Options for writeConcern	407
	The j (Journaling) Option for writeConcern	408
	Durability at a Cluster Level Using Read Concern	408
	Durability of Transactions Using a Write Concern	409
	What MongoDB Does Not Guarantee	410



Checking for Corruption	410
<hr/>	
<b>Part VI. Server Administration</b>	
<b>21. Setting Up MongoDB in Production.....</b>	<b>415</b>
Starting from the Command Line	415
File-Based Configuration	419
Stopping MongoDB	420
Security	421
Data Encryption	422
SSL Connections	423
Logging	423
<b>22. Monitoring MongoDB.....</b>	<b>425</b>
Monitoring Memory Usage	425
Introduction to Computer Memory	426
Tracking Memory Usage	426
Tracking Page Faults	427
I/O Wait	429
Calculating the Working Set	429
Some Working Set Examples	431
Tracking Performance	431
Tracking Free Space	433
Monitoring Replication	433
<b>23. Making Backups.....</b>	<b>437</b>
Backup Methods	437
Backing Up a Server	438
Filesystem Snapshot	438
Copying Data Files	442
Using mongodump	443
Specific Considerations for Replica Sets	446
Specific Considerations for Sharded Clusters	446
Backing Up and Restoring an Entire Cluster	447
Backing Up and Restoring a Single Shard	447
<b>24. Deploying MongoDB.....</b>	<b>449</b>
Designing the System	449
Choosing a Storage Medium	449



Recommended RAID Configurations	450
CPU	451
Operating System	451
Swap Space	452
Filesystem	452
Virtualization	453
Memory Overcommitting	453
Mystery Memory	453
Handling Network Disk I/O Issues	453
Using Non-Networked Disks	455
Configuring System Settings	455
Turning Off NUMA	455
Setting Readahead	457
Disabling Transparent Huge Pages (THP)	458
Choosing a Disk Scheduling Algorithm	458
Disabling Access Time Tracking	459
Modifying Limits	460
Configuring Your Network	461
System Housekeeping	462
Synchronizing Clocks	462
The OOM Killer	463
Turn Off Periodic Tasks	463
<b>A. Installing MongoDB.....</b>	<b>465</b>
<b>B. MongoDB Internals.....</b>	<b>469</b>
<b>Index.....</b>	<b>475</b>