

# Contents

	<b>Acknowledgments</b>	<b>vii</b>
	<b>Preface</b>	<b>ix</b>
<b>1.</b>	<b>Introduction</b>	<b>1</b>
	It Starts with a Question	2
	The Genres	3
	Onward and Upward	8
<b>2.</b>	<b>PostgreSQL</b>	<b>9</b>
	That's Post-greS-Q-L	9
	Day 1: Relations, CRUD, and Joins	10
	Day 2: Advanced Queries, Code, and Rules	21
	Day 3: Full Text and Multidimensions	36
	Wrap-Up	50
<b>3.</b>	<b>HBase</b>	<b>53</b>
	Introducing HBase	54
	Day 1: CRUD and Table Administration	55
	Day 2: Working with Big Data	67
	Day 3: Taking It to the Cloud	82
	Wrap-Up	88
<b>4.</b>	<b>MongoDB</b>	<b>93</b>
	Hu(mongo)us	93
	Day 1: CRUD and Nesting	94
	Day 2: Indexing, Aggregating, Mapreduce	110
	Day 3: Replica Sets, Sharding, GeoSpatial, and GridFS	124
	Wrap-Up	132
<b>5.</b>	<b>CouchDB</b>	<b>135</b>
	Relaxing on the Couch	135

Day 1: CRUD, Fauxton, and cURL Redux	137
Day 2: Creating and Querying Views	145
Day 3: Advanced Views, Changes API, and Replicating Data	158
Wrap-Up	174
<b>6. Neo4J</b>	<b>177</b>
Neo4j Is Whiteboard Friendly	177
Day 1: Graphs, Cypher, and CRUD	179
Day 2: REST, Indexes, and Algorithms	189
Day 3: Distributed High Availability	202
Wrap-Up	207
<b>7. DynamoDB</b>	<b>211</b>
DynamoDB: The “Big Easy” of NoSQL	211
Day 1: Let’s Go Shopping!	216
Day 2: Building a Streaming Data Pipeline	233
Day 3: Building an “Internet of Things” System Around DynamoDB	246
Wrap-Up	255
<b>8. Redis</b>	<b>259</b>
Data Structure Server Store	259
Day 1: CRUD and Datatypes	260
Day 2: Advanced Usage, Distribution	274
Day 3: Playing with Other Databases	289
Wrap-Up	303
<b>9. Wrapping Up</b>	<b>305</b>
Genres Redux	305
Making a Choice	309
Where Do We Go from Here?	309
<b>A1. Database Overview Tables</b>	<b>311</b>
<b>A2. The CAP Theorem</b>	<b>315</b>
Eventual Consistency	316
CAP in the Wild	317
The Latency Trade-Off	317
<b>Bibliography</b>	<b>319</b>
<b>Index</b>	<b>321</b>