

Seven Databases in Seven Weeks

A Guide to Modern Databases
and the NoSQL Movement

Eric Redmond
Jim R. Wilson

The Pragmatic Bookshelf

Dallas, Texas • Raleigh, North Carolina

Contents

Foreword	vii
Acknowledgments	ix
Preface	xi
1. Introduction	1
1.1 It Starts with a Question	1
1.2 The Genres	3
1.3 Onward and Upward	7
2. PostgreSQL	9
2.1 That's Post-greS-Q-L	9
2.2 Day 1: Relations, CRUD, and Joins	10
2.3 Day 2: Advanced Queries, Code, and Rules	21
2.4 Day 3: Full-Text and Multidimensions	35
2.5 Wrap-Up	48
3. Riak	51
3.1 Riak Loves the Web	51
3.2 Day 1: CRUD, Links, and MIMEs	52
3.3 Day 2: Mapreduce and Server Clusters	62
3.4 Day 3: Resolving Conflicts and Extending Riak	80
3.5 Wrap-Up	91
4. HBase	93
4.1 Introducing HBase	94
4.2 Day 1: CRUD and Table Administration	94
4.3 Day 2: Working with Big Data	106
4.4 Day 3: Taking It to the Cloud	122
4.5 Wrap-Up	131

5. MongoDB	135
5.1 Hu(mongo)us	135
5.2 Day 1: CRUD and Nesting	136
5.3 Day 2: Indexing, Grouping, Mapreduce	151
5.4 Day 3: Replica Sets, Sharding, GeoSpatial, and GridFS	165
5.5 Wrap-Up	174
6. CouchDB	177
6.1 Relaxing on the Couch	177
6.2 Day 1: CRUD, Futon, and cURL Redux	178
6.3 Day 2: Creating and Querying Views	186
6.4 Day 3: Advanced Views, Changes API, and Replicating Data	200
6.5 Wrap-Up	217
7. Neo4J	219
7.1 Neo4J Is Whiteboard Friendly	219
7.2 Day 1: Graphs, Groovy, and CRUD	220
7.3 Day 2: REST, Indexes, and Algorithms	238
7.4 Day 3: Distributed High Availability	250
7.5 Wrap-Up	258
8. Redis	261
8.1 Data Structure Server Store	261
8.2 Day 1: CRUD and Datatypes	262
8.3 Day 2: Advanced Usage, Distribution	275
8.4 Day 3: Playing with Other Databases	291
8.5 Wrap-Up	304
9. Wrapping Up	307
9.1 Genres Redux	307
9.2 Making a Choice	311
9.3 Where Do We Go from Here?	312
A1. Database Overview Tables	313
A2. The CAP Theorem	317
A2.1 Eventual Consistency	317
A2.2 CAP in the Wild	318
A2.3 The Latency Trade-Off	319
Bibliography	321
Index	323