



Riak

HANDBOOK

The hands-on guide to Riak
by Mathias Meyer

Riak Handbook

Mathias Meyer

Revision 27e1e7fb

Table of Contents

Introduction	8
Thank You	8
How to read the book.....	9
Feedback	9
Code.....	9
Changelog	9
CAP Theorem	11
The CAP Theorem is Not Absolute	12
Fine-Tuning CAP with Quorums	13
N, R, W, Quorums, Oh My!.....	13
How Quorums Affect CAP	14
A Word of CAP Wisdom.....	15
Further Reading	15
Eventual Consistency	15
Consistency in Quorum-Based Systems	16
Consistent Hashing	16
Sharding and Rehashing.....	16
A Better Way	17
Enter Consistent Hashing	17
Looking up an Object	19
Problems with Consistent Hashing	20
Dealing with Overload and Data Loss	21
Amazon's Dynamo	22
Basics.....	22
Virtual Nodes	22
Master-less Cluster	23
Quorum-based Replication	24
Read Repair and Hinted Handoff	24
Conflict Resolution using Vector Clocks	24
Conclusion	26
What is Riak?.....	27
Riak: Dynamo, And Then Some	27
Installation	28
Installing Riak using Binary Packages	28
Talking to Riak.....	29
Buckets	29
Fetching Objects	29
Creating Objects	30

Object Metadata	31
Custom Metadata	32
Linking Objects.....	33
Walking Links.....	34
Walking Nested Links	35
The Anatomy of a Bucket	36
List All Of The Keys.....	37
How Do I Delete All Keys in a Bucket?.....	38
How Do I Get the Number of All Keys in a Bucket?	39
Querying Data	39
MapReduce.....	40
MapReduce Basics	41
Mapping Tweet Attributes	41
Using Reduce to Count Tweets	42
Re-reducing for Great Good	43
Counting all Tweets.....	44
Chaining Reduce Phases	44
Parameterizing MapReduce Queries.....	46
Chaining Map Phases	48
MapReduce in a Riak Cluster.....	48
Efficiency of Buckets as Inputs.....	50
Key Filters.....	51
Using Riak's Built-in MapReduce Functions.....	53
Intermission: Riak's Configuration Files	54
Errors Running JavaScript MapReduce.....	55
Deploying Custom JavaScript Functions	56
Using Erlang for MapReduce	57
Writing Custom Erlang MapReduce Functions	58
On Full-Bucket MapReduce and Key-Filters Performance	61
Querying Data, For Real.....	61
Riak Search	62
Enabling Riak Search	62
Indexing Data.....	62
Indexing from the Command-Line.....	63
The Anatomy of a Riak Search Document.....	63
Querying from the Command-Line	64
Other Command-Line Features	64
The Riak Search Document Schema	64
Analyzers	65
Writing Custom Analyzers	66

Other Schema Options.....	69
An Example Schema.....	70
Setting the Schema.....	72
Indexing Data from Riak.....	72
Using the Solr Interface.....	74
Paginating Search Results.....	75
Sorting Search Results.....	76
Search Operators.....	76
Summary of Solr API Search Options.....	79
Summary of the Solr Query Operators.....	80
Indexing Documents using the Solr API.....	81
Deleting Documents using the Solr API.....	82
Using Riak's MapReduce with Riak Search.....	83
The Overhead of Indexing.....	83
Riak Secondary Indexes.....	84
Indexing Data with 2i.....	84
Querying Data with 2i.....	86
Using Riak 2i with MapReduce.....	87
Storing Multiple Index Values.....	87
Managing Object Associations: Links vs. 2i.....	88
How Does Riak 2i Compare to Riak Search?.....	89
Riak Search vs. Riak 2i vs. MapReduce.....	90
How Do I Index Data Already in Riak?.....	91
Using Pre- and Post-Commit Hooks.....	92
Validating Data.....	92
Enabling Pre-Commit Hooks.....	93
Pre-Commit Hooks in Erlang.....	94
Modifying Data in Pre-Commit Hooks.....	95
Accessing Riak Objects in Commit Hooks.....	97
Enabling Post-Commit Hooks.....	100
Deploying Custom Erlang Functions.....	100
Updating External Sources in Post-Commit Hooks.....	102
Riak in its Setting.....	102
Building a Cluster.....	102
Adding a Node to a Riak Cluster.....	103
Configuring a Riak Node.....	103
Joining a Cluster.....	104
Anatomy of a Riak Node.....	104
What Happens When a Node Joins a Cluster.....	105
Leaving a Cluster.....	105

Eventually Consistent Riak	106
Handling Consistency.....	106
Writing with a Non-Default Quorum.....	106
Durable Writes	107
Primary Writes	108
Tuning Default-Replication and Quorum Per Bucket.....	108
Choosing the Right N Value	110
Reading with a Non-Default Quorum.....	110
Read-Repair.....	111
Modeling Data for Eventual Consistency	111
Choosing the Right Data Structures.....	112
Conflicts in Riak.....	115
Siblings.....	116
Reconciling Conflicts.....	117
Modeling Counters and Other Data Structures	118
Problems with Timestamps for Conflict Resolution.....	119
Strategies for Reconciling Conflicts	123
Reads Before Writes	124
Merging Strategies	124
Sibling Explosion.....	124
Building a Timeline with Riak	125
Multi-User Timelines.....	128
Avoiding Infinite Growth.....	129
Intermission: How to Fetch Multiple Objects in one Request.....	129
Intermission: Paginating Using MapReduce	130
Handling Failure	131
Operating Riak.....	132
Choosing a Ring Size	132
Protocol Buffers vs. HTTP	133
Storage Backends.....	133
Innystore	134
Bitcask.....	134
LevelDB.....	135
Load-Balancing Riak	136
Placing Riak Nodes across a Network	138
Monitoring Riak.....	140
Request Times	141
Number of Requests	142
Read Repairs, Object Size, Siblings.....	143
Monitoring 2i	144

Miscellany	144
Monitoring Reference	144
Managing a Riak Cluster with Riak Control.....	147
Enabling Riak Control	147
Intermission: Generating an SSL Certificate	148
Riak Control Cluster Overview	149
Managing Nodes with Riak Control	150
Managing the Ring with Riak Control	151
To Be Continued.....	152
When To Riak?	152
Riak Use Cases in Detail.....	153
Using Riak for File Storage	153
File Storage Access Patterns	154
Object Size.....	154
Storing Large Files in Riak	155
Riak Cloud Storage	155
Using Riak to Store Logs.....	156
Modeling Log Records.....	157
Logging Access Patterns	157
Indexing Log Data for Efficient Access	158
Secondary Index Ranges as Key Filter Replacement	159
Searching Logs	160
Riak for Log Storage in the Wild.....	161
Deleting Historical Data	161
What about Analytics?	162
Session Storage	162
Modeling Session Data	163
Session Storage Access Patterns.....	164
Bringing Session Data Closer to Users	164
URL Shortener	164
URL Shortening Access Patterns	165
Modeling Data.....	165
Riak URL Shortening in the Wild.....	165
Where to go from here.....	165