

MongoDB

Preface

1. How This Book Is Organized
 1. Getting Started with MongoDB
 2. Developing with MongoDB
 3. Replication
 4. Sharding
 5. Application Administration
 6. Server Administration
 7. Appendixes
2. Conventions Used in This Book
3. Using Code Examples
4. O'Reilly Online Learning
5. How to Contact Us

I. Introduction to MongoDB

1. Introduction

1. Ease of Use
2. Designed to Scale
3. Rich with Features...
4. ...Without Sacrificing Speed
5. The Philosophy

2. Getting Started

1. Documents
2. Collections
 1. Dynamic Schemas
 2. Naming
3. Databases
4. Getting and Starting MongoDB
5. Introduction to the MongoDB Shell
 1. Running the Shell
 2. A MongoDB Client
 3. Basic Operations with the Shell
6. Data Types
 1. Basic Data Types
 2. Dates
 3. Arrays
 4. Embedded Documents
 5. _id and ObjectIds
7. Using the MongoDB Shell
 1. Tips for Using the Shell
 2. Running Scripts with the Shell
 3. Creating a .mongorc.js
 4. Customizing Your Prompt
 5. Editing Complex Variables
 6. Inconvenient Collection Names

3. Creating, Updating, and Deleting Documents

1. Inserting Documents
 1. insertMany
 2. Insert Validation
 3. insert
2. Removing Documents
 1. drop
3. Updating Documents
 1. Document Replacement
 2. Using Update Operators
 3. Upserts
 4. Updating Multiple Documents
 5. Returning Updated Documents

4. Querying

1. Introduction to find
 1. Specifying Which Keys to Return
 2. Limitations
2. Query Criteria
 1. Query Conditionals
 2. OR Queries
 3. \$not
3. Type-Specific Queries
 1. null
 2. Regular Expressions
 3. Querying Arrays
 4. Querying on Embedded Documents
4. \$where Queries
5. Cursors
 1. Limits, Skips, and Sorts
 2. Avoiding Large Skips
 3. Immortal Cursors

II. Designing Your Application

5. Indexes

1. Introduction to Indexes
 1. Creating an Index
 2. Introduction to Compound Indexes
 3. How MongoDB Selects an Index
 4. Using Compound Indexes
 5. How \$ Operators Use Indexes
 6. Indexing Objects and Arrays
 7. Index Cardinality
2. explain Output
3. When Not to Index
4. Types of Indexes
 1. Unique Indexes
 2. Partial Indexes

5. [Index Administration](#)
 1. [Identifying Indexes](#)
 2. [Changing Indexes](#)
6. [Special Index and Collection Types](#)
 1. [Geospatial Indexes](#)
 1. [Types of Geospatial Queries](#)
 2. [Using Geospatial Indexes](#)
 3. [Compound Geospatial Indexes](#)
 4. [2d Indexes](#)
 2. [Indexes for Full Text Search](#)
 1. [Creating a Text Index](#)
 2. [Text Search](#)
 3. [Optimizing Full-Text Search](#)
 4. [Searching in Other Languages](#)
 3. [Capped Collections](#)
 1. [Creating Capped Collections](#)
 2. [Tailable Cursors](#)
 4. [Time-To-Live Indexes](#)
 5. [Storing Files with GridFS](#)
 1. [Getting Started with GridFS: mongofiles](#)
 2. [Working with GridFS from the MongoDB Drivers](#)
 3. [Under the Hood](#)
7. [Introduction to the Aggregation Framework](#)
 1. [Pipelines, Stages, and Tunables](#)
 2. [Getting Started with Stages: Familiar Operations](#)
 3. [Expressions](#)
 4. [\\$project](#)
 5. [\\$unwind](#)
 6. [Array Expressions](#)
 7. [Accumulators](#)
 1. [Using Accumulators in Project Stages](#)
 8. [Introduction to Grouping](#)
 1. [The _id Field in Group Stages](#)
 2. [Group Versus Project](#)
 9. [Writing Aggregation Pipeline Results to a Collection](#)
8. [Transactions](#)
 1. [Introduction to Transactions](#)
 1. [A Definition of ACID](#)
 2. [How to Use Transactions](#)
 3. [Tuning Transaction Limits for Your Application](#)
 1. [Timing and Oplog Size Limits](#)
9. [Application Design](#)
 1. [Schema Design Considerations](#)
 1. [Schema Design Patterns](#)
 2. [Normalization Versus Denormalization](#)
 1. [Examples of Data Representations](#)

2. [Cardinality](#)
3. [Friends, Followers, and Other Inconveniences](#)
3. [Optimizations for Data Manipulation](#)
 1. [Removing Old Data](#)
4. [Planning Out Databases and Collections](#)
5. [Managing Consistency](#)
6. [Migrating Schemas](#)
7. [Managing Schemas](#)
8. [When Not to Use MongoDB](#)

[III. Replication](#)

[10. Setting Up a Replica Set](#)

1. [Introduction to Replication](#)
2. [Setting Up a Replica Set, Part 1](#)
3. [Networking Considerations](#)
4. [Security Considerations](#)
5. [Setting Up a Replica Set, Part 2](#)
6. [Observing Replication](#)
7. [Changing Your Replica Set Configuration](#)
8. [How to Design a Set](#)
 1. [How Elections Work](#)
9. [Member Configuration Options](#)
 1. [Priority](#)
 2. [Hidden Members](#)
 3. [Election Arbiters](#)
 4. [Building Indexes](#)

[11. Components of a Replica Set](#)

1. [Syncing](#)
 1. [Initial Sync](#)
 2. [Replication](#)
 3. [Handling Staleness](#)
2. [Heartbeats](#)
 1. [Member States](#)
3. [Elections](#)
4. [Rollbacks](#)
 1. [When Rollbacks Fail](#)

[12. Connecting to a Replica Set from Your Application](#)

1. [Client-to-Replica Set Connection Behavior](#)
2. [Waiting for Replication on Writes](#)
 1. [Other Options for “w”](#)
3. [Custom Replication Guarantees](#)
 1. [Guaranteeing One Server per Data Center](#)
 2. [Guaranteeing a Majority of Nonhidden Members](#)
 3. [Creating Other Guarantees](#)
4. [Sending Reads to Secondaries](#)
 1. [Consistency Considerations](#)
 2. [Load Considerations](#)

3. [Reasons to Read from Secondaries](#)

[13. Administration](#)

1. [Starting Members in Standalone Mode](#)
2. [Replica Set Configuration](#)
 1. [Creating a Replica Set](#)
 2. [Changing Set Members](#)
 3. [Creating Larger Sets](#)
 4. [Forcing Reconfiguration](#)
3. [Manipulating Member State](#)
 1. [Turning Primaries into Secondaries](#)
 2. [Preventing Elections](#)
4. [Monitoring Replication](#)
 1. [Getting the Status](#)
 2. [Visualizing the Replication Graph](#)
 3. [Replication Loops](#)
 4. [Disabling Chaining](#)
 5. [Calculating Lag](#)
 6. [Resizing the Oplog](#)
 7. [Building Indexes](#)
 8. [Replication on a Budget](#)

[IV. Sharding](#)

[14. Introduction to Sharding](#)

1. [What Is Sharding?](#)
 1. [Understanding the Components of a Cluster](#)
2. [Sharding on a Single-Machine Cluster](#)

[15. Configuring Sharding](#)

1. [When to Shard](#)
2. [Starting the Servers](#)
 1. [Config Servers](#)
 2. [The mongos Processes](#)
 3. [Adding a Shard from a Replica Set](#)
 4. [Adding Capacity](#)
 5. [Sharding Data](#)
3. [How MongoDB Tracks Cluster Data](#)
 1. [Chunk Ranges](#)
 2. [Splitting Chunks](#)
4. [The Balancer](#)
5. [Collations](#)
6. [Change Streams](#)

[16. Choosing a Shard Key](#)

1. [Taking Stock of Your Usage](#)
2. [Picturing Distributions](#)
 1. [Ascending Shard Keys](#)
 2. [Randomly Distributed Shard Keys](#)
 3. [Location-Based Shard Keys](#)
3. [Shard Key Strategies](#)

1. [Hashed Shard Key](#)
2. [Hashed Shard Keys for GridFS](#)
3. [The Firehose Strategy](#)
4. [Multi-Hotspot](#)
4. [Shard Key Rules and Guidelines](#)
 1. [Shard Key Limitations](#)
 2. [Shard Key Cardinality](#)
5. [Controlling Data Distribution](#)
 1. [Using a Cluster for Multiple Databases and Collections](#)
 2. [Manual Sharding](#)

[17. Sharding Administration](#)

1. [Seeing the Current State](#)
 1. [Getting a Summary with sh.status\(\)](#)
 2. [Seeing Configuration Information](#)
2. [Tracking Network Connections](#)
 1. [Getting Connection Statistics](#)
 2. [Limiting the Number of Connections](#)
3. [Server Administration](#)
 1. [Adding Servers](#)
 2. [Changing Servers in a Shard](#)
 3. [Removing a Shard](#)
4. [Balancing Data](#)
 1. [The Balancer](#)
 2. [Changing Chunk Size](#)
 3. [Moving Chunks](#)
 4. [Jumbo Chunks](#)
 5. [Refreshing Configurations](#)

[V. Application Administration](#)

[18. Seeing What Your Application Is Doing](#)

1. [Seeing the Current Operations](#)
 1. [Finding Problematic Operations](#)
 2. [Killing Operations](#)
 3. [False Positives](#)
 4. [Preventing Phantom Operations](#)
2. [Using the System Profiler](#)
3. [Calculating Sizes](#)
 1. [Documents](#)
 2. [Collections](#)
 3. [Databases](#)
4. [Using mongotop and mongostat](#)

[19. An Introduction to MongoDB Security](#)

1. [MongoDB Authentication and Authorization](#)
 1. [Authentication Mechanisms](#)
 2. [Authorization](#)
 3. [Using x.509 Certificates to Authenticate Both Members and Clients](#)
2. [A Tutorial on MongoDB Authentication and Transport Layer Encryption](#)

1. [Establish a CA](#)
2. [Generate and Sign Member Certificates](#)
3. [Generate and Sign Client Certificates](#)
4. [Bring Up the Replica Set Without Authentication and Authorization Enabled](#)
5. [Create the Admin User](#)
6. [Restart the Replica Set with Authentication and Authorization Enabled](#)

[20. Durability](#)

1. [Durability at the Member Level Through Journaling](#)
2. [Durability at the Cluster Level Using Write Concern](#)
 1. [The w and wtimeout Options for writeConcern](#)
 2. [The j \(Journaling\) Option for writeConcern](#)
3. [Durability at a Cluster Level Using Read Concern](#)
4. [Durability of Transactions Using a Write Concern](#)
5. [What MongoDB Does Not Guarantee](#)
6. [Checking for Corruption](#)

[VI. Server Administration](#)

[21. Setting Up MongoDB in Production](#)

1. [Starting from the Command Line](#)
 1. [File-Based Configuration](#)
2. [Stopping MongoDB](#)
3. [Security](#)
 1. [Data Encryption](#)
 2. [SSL Connections](#)
4. [Logging](#)

[22. Monitoring MongoDB](#)

1. [Monitoring Memory Usage](#)
 1. [Introduction to Computer Memory](#)
 2. [Tracking Memory Usage](#)
 3. [Tracking Page Faults](#)
 4. [I/O Wait](#)
2. [Calculating the Working Set](#)
 1. [Some Working Set Examples](#)
3. [Tracking Performance](#)
4. [Tracking Free Space](#)
5. [Monitoring Replication](#)

[23. Making Backups](#)

1. [Backup Methods](#)
2. [Backing Up a Server](#)
 1. [Filesystem Snapshot](#)
 2. [Copying Data Files](#)
 3. [Using mongodump](#)
3. [Specific Considerations for Replica Sets](#)
4. [Specific Considerations for Sharded Clusters](#)
 1. [Backing Up and Restoring an Entire Cluster](#)
 2. [Backing Up and Restoring a Single Shard](#)

[24. Deploying MongoDB](#)

5. [Designing the System](#)
 1. [Choosing a Storage Medium](#)
 2. [Recommended RAID Configurations](#)
 3. [CPU](#)
 4. [Operating System](#)
 5. [Swap Space](#)
 6. [Filesystem](#)
6. [Virtualization](#)
 1. [Memory Overcommitting](#)
 2. [Mystery Memory](#)
 3. [Handling Network Disk I/O Issues](#)
 4. [Using Non-Networked Disks](#)
7. [Configuring System Settings](#)
 1. [Turning Off NUMA](#)
 2. [Setting Readahead](#)
 3. [Disabling Transparent Huge Pages \(THP\)](#)
 4. [Choosing a Disk Scheduling Algorithm](#)
 5. [Disabling Access Time Tracking](#)
 6. [Modifying Limits](#)
8. [Configuring Your Network](#)
9. [System Housekeeping](#)
 1. [Synchronizing Clocks](#)
 2. [The OOM Killer](#)
 3. [Turn Off Periodic Tasks](#)