

MongoDB

Day 1

1. Databases
2. Getting and Starting MongoDB
3. Introduction to the MongoDB Shell
 1. Running the Shell
 2. A MongoDB Client
 3. Basic Operations with the Shell
4. Data Types
 1. Basic Data Types
 2. Dates
 3. Arrays
 4. Embedded Documents
 5. `_id` and `ObjectId`s
5. 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
6. Creating, Updating, and Deleting Documents
7. Inserting Documents
 1. `insertMany`
 2. Insert Validation
 3. `insert`
8. Removing Documents
9. Updating Documents
 1. Document Replacement
 2. Using Update Operators
 3. Upserts
 4. Updating Multiple Documents
 5. Returning Updated Documents

Day 2

10. 4. Querying
 1. Introduction to `find`
11. Specifying Which Keys to Return
12. Limitations
13. Query Criteria
 1. Query Conditionals
 2. OR Queries
 3. `$not`
14. Type-Specific Queries

1. null
 2. Regular Expressions
 3. Querying Arrays
 4. Querying on Embedded Documents
15. \$where Queries
16. Cursors
 1. Limits, Skips, and Sorts
 2. Avoiding Large Skips
 3. Immortal Cursors
17. II. Designing Your Application 5. Indexes
18. 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
19. explain Output
20. When Not to Index
21. Types of Indexes
 1. Unique Indexes
 2. Partial Indexes
22. Index Administration
23. Identifying Indexes
24. Changing Indexes

Day 3

25. 6. Special Index and Collection Types
26. Geospatial Indexes
 1. Types of Geospatial Queries
 2. Using Geospatial Indexes
 3. Compound Geospatial Indexes
 4. 2d Indexes
27. Indexes for Full Text Search
 1. Creating a Text Index
 2. Text Search
 3. Optimizing Full-Text Search
 4. Searching in Other Languages
28. Capped Collections
 1. Creating Capped Collections
 2. Tailable Cursors
29. Time-To-Live Indexes
30. Storing Files with GridFS
 1. Getting Started with GridFS: mongofiles
 2. Working with GridFS from the MongoDB Drivers

- 3. Under the Hood
- 31. Accumulators
 - 1. Using Accumulators in Project Stages
- 32. Introduction to Grouping
 - 1. The `_id` Field in Group Stages
 - 2. Group Versus Project
- 33. Writing Aggregation Pipeline Results to a Collection
- 34. 8. Transactions
 - 1. Introduction to Transactions
- 35. A Definition of ACID
- 36. How to Use Transactions
- 37. Tuning Transaction Limits for Your Application
 - 1. Timing and Oplog Size Limits
- 38. Application Design
 - 1. Schema Design Considerations
 - 1. Schema Design Patterns
 - 2. Normalization Versus Denormalization

Day 4

- 39. Examples of Data Representations
- 40. Cardinality
- 41. Friends, Followers, and Other Inconveniences
- 3. Optimizations for Data Manipulation
- 42. Removing Old Data
- 43. Planning Out Databases and Collections
- 44. Managing Consistency
- 45. Migrating Schemas
- 46. Managing Schemas
- 47. When Not to Use MongoDB
- 48. III. Replication
 - 10. Setting Up a Replica Set
- 49. Introduction to Replication
- 50. Setting Up a Replica Set, Part 1
- 51. Networking Considerations
- 52. Security Considerations
- 53. Setting Up a Replica Set, Part 2
- 54. Observing Replication
- 55. Changing Your Replica Set Configuration
- 56. How to Design a Set

Day 5

- 57. How Elections Work
 - 9. Member Configuration Options
- 58. Priority
- 59. Hidden Members
- 60. Election Arbiters

- 61. Building Indexes
- 62. Components of a Replica Set 1. Syncing
- 63. Initial Sync
- 64. Replication
- 65. Handling Staleness
- 66. Heartbeats
 - 1. Member States
- 67. 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
- 68. Configuring Your Network
- 69. System Housekeeping
 - 1. Synchronizing Clocks
 - 2. The OOM Killer
 - 3. Turn Off Periodic Tasks