



Mastering MongoDB 7.0 : Achieve Data Excellence by Unlocking the Full Potential of MongoDB /

Aleksendric, Marko,
author

Monografía

Mastering MongoDB 7.0 explores the latest version of MongoDB, an exceptional NoSQL database solution that aligns with the needs of modern web applications. This book starts with an informative overview of MongoDB's architecture and developer tools, guiding you through the process of connecting to databases seamlessly. This MongoDB book explores advanced queries in detail, including aggregation pipelines and multi-document ACID transactions. It delves into the capabilities of the MongoDB Atlas developer data platform and the latest features, such as Atlas Vector Search, and their role in AI applications, enabling developers to build applications with the scalability and performance that today's organizations need. It also covers the creation of resilient search functionality using MongoDB Atlas Search. Mastering MongoDB 7.0's deep coverage of advanced techniques encompasses everything from role-based access control (RBAC) to user management, auditing practices, and encryption across data, network, and storage layers. By the end of this book, you'll have developed the skills necessary to create efficient, secure, and high-performing applications using MongoDB. You'll have the confidence to undertake complex queries, integrate robust applications, and ensure data security to overcome modern data challenges

<https://rebiunoda.pro.baratznet.cloud:38443/OpacDiscovery/public/catalog/detail/b2FpOmNlbGVicmF0aW9uOmVzLmJhcmF0ei5yZW4vMzUyOTM3MDg>

Título: Mastering MongoDB 7.0 Achieve Data Excellence by Unlocking the Full Potential of MongoDB Marko Aleksendric [and six others]

Edición: 4th ed

Editorial: Birmingham, England Packt Publishing [2024] 2024

Descripción física: 1 online resource (447 pages)

Nota general: Includes index

Contenido: Cover -- Title Page -- Copyright and Credits -- Contributors -- Note from Author -- Table of Contents -- Preface -- Chapter 1: Introduction to MongoDB -- Who uses MongoDB? -- Why developers love MongoDB? -- Efficiency of the inherent complexity of MongoDB databases -- How is this possible? -- Summary -- Chapter 2: The MongoDB Architecture -- Replication vs sharding -- Replication -- Replica set elections -- Replica set oplog -- Replica set deployment architectures -- Write concern -- Read preference -- Read concern -- Replication methods --

Sharding -- Why do you need sharding? -- Key elements of a sharded cluster -- Microsharding -- Advantages of sharding -- Data distribution -- Chunks -- Chunk administration -- Querying sharded data -- Sharding methods -- New sharded cluster features in MongoDB 7.0 -- Shard key advisor commands -- AutoMerger -- Command support without shard keys -- Summary -- Chapter 3: Developer Tools -- Technical requirements -- Introduction to developer tools -- Categories of developer tools -- MongoDB Shell -- Installation -- Connecting to a MongoDB deployment using mongosh -- mongosh vs legacy shell -- Key features of mongosh -- Use cases for mongosh -- Best practices for using mongosh -- MongoDB CLI -- Installation -- Configuration -- Working with MongoDB CLI -- mongocli auth -- Key features of MongoDB CLI -- Use cases for MongoDB CLI -- Best practices for MongoDB CLI -- MongoDB Compass -- Installation -- Configuration -- Visually explore data using Compass -- Key features of MongoDB Compass -- Use cases for MongoDB Compass -- Best practices for MongoDB Compass -- MongoDB for VS Code -- Installation -- Explore data using Playground -- Use Terraform to create an Atlas cluster -- Key features -- Use cases for the MongoDB for VS Code extension -- Best practices for the MongoDB for VS Code extension -- Summary -- Chapter 4: Connecting to MongoDB -- Connecting methods -- Connecting using Ruby -- Connecting using Python -- Connecting using PHP -- Connecting using Node.js -- Connecting using TypeScript -- Summary -- Chapter 5: CRUD Operations and Basic Queries -- Technical requirements -- MongoDB CRUD operations -- CRUD using mongosh -- Scripting for mongosh -- Batch inserts using mongosh -- Batch operations using mongosh -- CRUD using the Ruby driver -- Connecting to a database -- Creating documents -- Reading data -- Chaining operations in find() -- Nested operations -- Batch operations -- Mongoid ODM -- CRUD using the Python driver -- Inserting documents -- Finding documents -- Updating documents -- Deleting documents -- Regular expressions -- Administration -- currentOp() and killOp() -- collMod -- Secure access to MongoDB -- Authentication and authorization -- MongoDB Stable API -- Summary -- Chapter 6: Schema Design and Data Modeling -- Schema design for relational databases -- Normal forms -- Schema design for MongoDB -- Data modeling in MongoDB -- Document structure -- Embedded data -- References -- Denormalization -- BSON and its data types -- Design considerations and best practices for MongoDB modeling -- Read-write ratio -- Design patterns and schema design -- Summary -- Chapter 7: Advanced Querying in MongoDB -- Introduction to the aggregation framework -- Benefits of aggregation in MongoDB -- Aggregation stages -- Query Techniques -- Logical and comparison operators -- Array querying and manipulation -- Array field projection techniques -- Indexes and query optimization -- Benefits of using indexes -- Types of index -- Geospatial features in MongoDB -- GeoJSON objects -- Geospatial indexes -- Geospatial operators -- Summary -- Chapter 8: Aggregation -- Technical requirements -- MongoDB aggregation framework -- Modularity and composability -- Stages -- Aggregation expressions -- Basic aggregation operators -- Bitwise operators in aggregation -- Array expression operators -- The "power" array operators: map and reduce -- Other types of aggregations -- Best practices -- Code modularity -- Query optimization -- Streaming and blocking stages of a pipeline -- Sharded cluster considerations -- Summary -- Chapter 9: Multi-Document ACID Transactions -- Why are transactions useful? -- ACID properties -- Atomicity -- Consistency -- Isolation -- Durability -- MongoDB implementation of ACID -- Read and write concerns in transactions -- Transaction limitations -- Best practices -- Summary -- Chapter 10: Index Optimization -- Introduction to indexes -- Creating an index -- Index types -- Single-field indexes and indexing embedded documents -- Compound indexes -- Multikey indexes -- Covered queries -- Hashed index -- Text indexes -- Atlas Search Indexes -- TTL indexes -- Partial indexes -- Geospatial indexes -- Wildcard indexes -- Hidden indexes -- Compound wildcard indexes -- Best practices for index optimization -- Summary -- Chapter 11: MongoDB Atlas: Powering the Future of Developer Data Platforms -- Understanding MongoDB Atlas as a database-as-a-service -- Hosted database -- Secure by default -- Monitoring and alerting -- Back up, restore, and archive -- Automation -- Atlas developer data platform -- Atlas Vector Search and its role in AI applications -- What does vector search entail? -- What is Atlas Vector Search? -- Semantic search and vectors -- KNN, ANN, and vector embeddings -- How to use Atlas Vector Search -- Atlas Search -- Atlas Application Services -- What is Atlas Application Services? -- Key features -- How to use Atlas Application Services -- Atlas Data API -- Security and authorization -- Use cases -- How to use the Atlas Data API -- Atlas Data Lake -- Data storage and optimization -- How to use Atlas Data Lake -- Atlas Data Federation -- Data processing -- Key features -- How to use Atlas Data Federation -- Atlas Stream Processing (preview) -- Operational mechanism -- Key features -- Setting up a stream processor -- Atlas SQL Interface -- How to use the Atlas SQL Interface -- MongoDB Atlas Charts -- How to use Atlas Charts -- Operational integrations: The Atlas Kubernetes Operator -- What is the Atlas Kubernetes Operator? -- How to use the Atlas Kubernetes Operator -- Atlas CLI -- How to use the Atlas CLI -- Capabilities of the Atlas CLI -- Summary -- Chapter 12: Monitoring and Backup in MongoDB -- Monitoring in MongoDB -- Monitoring clusters --

Why monitor in MongoDB? -- What should you monitor? -- Page fault -- Resident memory -- Virtual and mapped memory -- Working set -- Track free disk space -- Monitoring replication -- Oplog size -- Network monitoring -- Cursors and connections -- Document metrics -- Monitoring memory usage in WiredTiger -- Tracking page faults -- Tracking B-tree misses -- I/O wait -- Read and write queues -- Working set calculations -- Why is the working set crucial? -- Estimate the working set size -- Importance of indexes in the working set -- MongoDB reporting tools overview -- Utilities -- Commands -- Hosted monitoring tools overview -- Cluster backups -- MongoDB backup methods -- Cloud-based backup solutions -- On-premises backup solutions -- Using filesystem snapshots for MongoDB backups -- Consistency requirements -- Backups using mongodump -- Backing up with cp or rsync -- Backup and restore on EC2 -- Backup with queues -- Incremental backups -- Common mistakes and pitfalls in MongoDB monitoring and backup -- Summary -- Chapter 13: Introduction to Atlas Search -- MongoDB Atlas Search -- Feature overview -- Version requirements for Atlas Search features -- Atlas Search architecture Technical overview of Atlas Search indexes -- Index creation and management -- Atlas Search queries -- Atlas Search query operators and collectors -- Query execution in Atlas Search -- Relevance scoring in Atlas Search -- Atlas Search Nodes -- Atlas Search index management -- Search query analysis -- Troubleshooting Atlas Search errors -- Apache Lucene -- Apache Software Foundation -- Atlas Search and integration with Apache Lucene -- Summary -- Chapter 14: Integrating Applications with MongoDB -- Technical requirements -- Integrating applications in MongoDB -- The MongoDB Kubernetes Operator -- Installation -- Deploy and configure MongoDB using the Kubernetes Operator -- The Atlas Kubernetes Operator -- The Enterprise Kubernetes Operator -- Integrating Terraform with MongoDB -- Installation -- Configuring a MongoDB cluster on Atlas using Terraform -- Using Vercel with MongoDB -- Building a web application with Vercel and MongoDB -- Configuring MongoDB Atlas with Vercel -- Disconnecting the Vercel integration -- Integrating Datadog with MongoDB -- Configuring Datadog -- Integrating MongoDB Atlas with Datadog -- Integrating Prometheus with MongoDB -- Configuring Prometheus -- Integrating MongoDB Atlas with Prometheus -- Integrating webhooks with MongoDB -- Configuring webhooks with MongoDB -- Integrating MongoDB with alerting applications -- Atlas integrations for webhooks -- Microsoft Teams integration -- PagerDuty integration -- Slack integration -- Splunk On-Call integration -- Opsgenie integration -- Integrating MongoDB with Confluent -- Integrating MongoDB with Postman -- Summary -- Chapter 15: Security -- Authentication methods -- SCRAM -- x.509 -- LDAP -- Kerberos -- Role-based access control (RBAC) -- Knowing about resources, actions, and roles -- Role hierarchy -- Managing roles -- Access restrictions -- Summary -- Chapter 16: Auditing Auditing and logging in MongoDB

ISBN: 1-83546-319-3

Materia Título preferido: MongoDB

Materia: Open source software Object-oriented databases Database management- Software

Enlace a formato físico adicional: 9781835460474

Baratz Innovación Documental

- Gran Vía, 59 28013 Madrid
- (+34) 91 456 03 60
- informa@baratz.es