



MongoDB Fundamentals

Tecnologias de Informação - Data & Analytics

Live Training (também disponível em presencial)

- **Localidade:** Lisboa
 - **Data:** 18 Sep 2023
 - **Preço:** 1080 € (Os valores apresentados não incluem IVA. Oferta de IVA a particulares e estudantes.)
 - **Horário:** Laboral das 09h30 - 16h30
 - **Nível:**
 - **Duração:** 18h
-

Sobre o curso

In this course you will learn how to set up your database and start exploring different ways to search, create, and analyze your data with MongoDB. We will cover database performance basics, and discover how to get started with creating applications and visualizing your data.

You will get your hands on all the basics, including querying, computing, connecting to, storing, indexing and analyzing your data.

Destinatários

- This course targets professional developers who are looking for in-depth knowledge for developing, administering, securing and scaling MongoDB-based applications.
-

Objetivos

- Develop highly scalable and cost-efficient applications with MongoDB
- Build efficient MongoDB data models for enterprise-scale applications
- Enhance performance with indexing and sharding
- Manage consistency with write concern and read preference
- Process data and compute results with aggregation pipelines and MapReduce

Pré-requisitos

- Participants of this course need to have a good understanding Linux and Python
-

Programa

- Interacting with MongoDB
- Installing MongoDB
- Exploring the MongoDB Data and Query Model
- Configuring Replication and Data Consistency
- Tuning Performance with Indexes
- Utilizing the Aggregation Framework

Interacting with MongoDB

- Justifying non-relational data stores
- Listing the categories of NoSQL Data Stores
- Defining the MongoDB architecture
- Reviewing the JSON standard
- Enumerating the MongoDB data types
- Querying at the Command Line Interface
- Running Scripts
- Manipulating MongoDB data from Python and Java

Installing MongoDB

- Surveying Cluster Topology
- Creating a standalone mongod
- Achieving redundancy with replica sets
- Replicating data with the oplog
- Identifying storage engines
- Improving data safety with journaling
- Investigating record allocation strategies
- Weighing the advantages and disadvantages of GridFS
- Leveraging the mongofiles utility tool
- Compressing data for economy of drive space

Exploring the MongoDB Data and Query Model

- Storing data in Databases, collections and documents
- Inserting and updating Documents
- Running scripts from the Mongo Shell
- Restricting find() with a query document
- Specifying keys for queries to return
- Constraining the number of results with limit and skip
- Express relationships using references
- Embedding sub-documents

Configuring Replication and Data Consistency

- Configuring Read and Write Consistency
- Reading from primary and secondary nodes
- Setting the read preference level
- Selecting the Appropriate Write Concern
- Writing with a timeout
- Applying tags to replicas
- Adding logical names to write concerns
- Reading from a tagged mongod

Tuning Performance with Indexes

- Investigating a query with the explain() function
- Getting statistics from explain()
- Choosing an index with the hint() function
- Designing a better index
- Opting out of using an index
- Profiling to identify slow queries
- Sorting and finding ranges with an index
- Guaranteeing uniqueness with indexes
- Skipping missing fields with a sparse index
- Aging data out of a capped collection
- Searching for text fragments with \$text
- Deleting documents automatically using time to live indexes
- Permitting queries by location

Utilizing the Aggregation Framework

- Leveraging the different modes of the Aggregation Framework
- Invoking single-purpose aggregation functions
- Performing operations with the Aggregation Pipeline framework
- Executing MapReduce on MongoDB