



## MongoDB for Software Engineers

24 hours

### Course Overview:

MongoDB is a leading document database that changes the world developers know so far: no schema constraints, instant object persistency, rapid high availability, amazing performance and scale out support for the cloud generation.

This course provides a developer level introduction along with more advanced and useful features. The course will include hands-on practice with MongoDB and Java, creating a complete system from data design to implementation using Java and integration with BI and dashboard tools.

Note: This course can be adjusted to additional programming languages (except for Java) and other target audiences including DBA and DevOps based on customer request.

### Who Should Attend:

Software development managers, CTO, software architects, system architects, data architects and developers.

### Required Skills:

Software development and database design, experience with Linux and access to a Linux machine.

### Course Contents:

#### Module 1: Introduction to NoSQL

- CAP Theorem
- What are the main concerns with RDBMS (SQL Server, Oracle, MySQL)
- Key-value stores
- Column Family stores
- Document DBs
- Map Reduce

#### Module 2: Introduction to MongoDB

- MongoDB product design and architecture
- MongoDB installation
- The Mongo Shell
- Basic Operations
- Lab: MongoDB installation and basic operations

#### Module 3: Data Model Design

- Documents and collections core concepts
- Data Model migration from relational DB to document store
- Data model best practices
- Lab: Data model design for the case

#### Module 4: CRUD

- Select
- Update

- Insert
- Delete
- Atomic Transactions
- Bulk Operations
- Lab: System implementation based on Java and MongoDB

**Module 5: Tuning your code**

- Indexing
- Query profiling
- The query optimizer
- Explain
- Lab: Profile and tune queries

**Module 6: Backup, Security and Monitoring**

- Dump
- Physical files backup
- Authentication
- Server status and system monitoring
- Lab: Understanding server performance metrics

**Module 7: Scale and High Availability**

- Data replication with replica sets
- Load distribution with sharding
- Performance best practices: leveraging scale to meet performance needs

**Module 8: Advanced Querying and Leveraging Scale**

- The Aggregation Framework
- Map Reduce
- Lab: Advanced querying using MongoDB