

BIG DATA AND NOSQL FOR DEVELOPERS TRAINING

(2 Day)

This course provides application developers with technical overview of Big Data as well as NoSQL (Not Only SQL) database systems. Effective use of NoSQL systems and understanding the appropriate ways of handling Big Data leads to the creation of the next-generation of high-performance and robust solutions.

Objectives

This intensive training course aims at making students cognisant of NoSQL systems capabilities and how they can be leveraged.

The course is supplemented by hands-on labs that help attendees reinforce their theoretical knowledge of the learned material and make them confident in applying the acquired knowledge in practice.

Topics

- Defining Big Data
- Big Data Stores Overview
- NoSQL
- Big Data Business Intelligence and Analytics
- Google App Engine Interfaces
- Working with MongoDB

Audience

Technical leads and application developers

Pre-requisites

Participants should be familiar with programming in Java and using Eclipse development environment

Course Outline

1. Defining Big Data

- Transforming Data into Business Information
- Definition of Big Data
- Challenges Posed by Big Data
- The Cloud and Big Data
- The Business Value of Big Data
- Big Data: Hype or Reality?

2. Big Data Systems Overview

- Limitations of Relational Databases
- NoSQL Database Systems
- The CAP theorem
- Limitations of NoSQL Databases
- Big Data Sharding
- Amazon S3
- Amazon S3 Security
- Data Lifecycle Management with Amazon S3
- Amazon S3 Cost Monitoring
- Google BigTable
- BigTable Design
- Google App Engine
- App Engine Billing
- Hadoop
- Hadoop Core Components
- Hadoop Distributed File System
- HBase
- HBase Design
- Cassandra
- Neo4J
- MemcacheDB
- MongoDB
- MongoDB Operational Intelligence
- MongoDB Use Cases

3. Big Data Business Intelligence and Analytics

- Comparison with Other Systems
- NoSQL Data Querying and Processing
- MapReduce Framework
- Analyzing Big Data with Hadoop
- Making things simpler with Pig Latin
- Example of a Pig Script in Batch Mode
- Amazon Elastic MapReduce
- Business Analytics with Hive
- The UnQL Specification

4. Working with Google App Engine

- Runtime Environments
- Development Environment
- Big Data in Google App Engine
- App Engine Datastore
- Google Cloud SQL
- Google Cloud Storage
- Blobstore
- Java Data Store API
- App Engine Services
- App Identity
- Memcache
- OAuth
- Task Queues
- URL Fetch
- XMPP

5. Working with MongoDB

- Drivers and Client Libraries
- MongoDB Data Model
- Administration Console
- Security and Authentication
- Data import and Export
- Managing MongoDB Lifecycles
- Read Operations
- Cursors
- Write Operations
- Querying, Limiting, Sorting and Aggregating data
- MongoDB QL
- Optimizing Queries with Indexes

Objectives

This intensive hands-on training course aims at making students proficient in developing with MongoDB using Java.

The course is supplemented by hands-on labs that help attendees reinforce their theoretical knowledge of the learned material and make them confident in applying the acquired knowledge in practice.

Topics

- Overview of "NoSQL" landscape
- MongoDB developer value proposition
- MongoDB architecture
- MongoDB installation and configuration
- Essential system administration
- Developing solutions in Java

Audience

Technical leads and application developers

Prerequisites

Participants should be familiar with Java programming and using Eclipse development environment

Duration

2 Days

CHAPTER 1. MONGODB OVERVIEW

- Defining "NoSQL"
- Introducing MongoDB
- MongoDB architecture
- Drivers and Client Libraries

CHAPTER 2. MONGODB DATA MODEL

- Databases
- Collections
- Documents
- BSON data format
- Indexes

CHAPTER 3. MONGODB ADMINISTRATION

- Administration Console
- Security and Authentication
- Basic MongoDB commands
- Data import and export
- Managing MongoDB lifecycles

CHAPTER 4. DATA STORAGE

- Data sharding concepts
- Chunking with GridFS

CHAPTER 5. READ AND WRITE OPERATIONS

- Read Operations
- Cursors
- Write Operations

CHAPTER 6. QUERYING, LIMITING, SORTING AND AGGREGATING DATA

- Querying
- Aggregations operations
- Map-reduce

CHAPTER 7. INDEXES

- Index Types
- Index Properties
- Index Creation

CHAPTER 8. ACCESSING MONGODB WITH JAVA

