

Introduction to NoSQL Training

C123-103

Learn and experience Big Data Solutions including NoSQL DBs: Document Based DB (Mongo DB), Big Table Databases (Apache Cassandra DB) Column Based DB Apache HBase In Memory Key Data store (Redis) AND In Memory Data Grids (Apache Ignite). Choose best suited technologies meeting your needs

This training introduces the popular NOSQL current solutions and provides basic hands on experience on each of the solutions themselves.

AUDIENCE

System Administrators
Operations
Database Administrators
Support
DevOps
Developers

KNOWLEDGE REQUIREMENTS

Relational Database SQL experience

LENGTH

3 Days

BONUS

Hands-on lab sessions

SYLLABUS

Course Introduction
Introduction to Big Data and NoSQL
Quick Introduction to Big Table Databases – Apache Cassandra DB
Quick Introduction to Document Databases – Mongo DB
Quick Introduction to Column Based Databases – Hadoop HBase DB
Introduction to Redis
Introduction to SPARK
Data Architect for Big Data
NOSQL Comparison
Summary

Customer can choose between any of the following:

Kafka, Elastic search, Redis, Apache Ignite, Apache Cassandra, Mongo DB, Gigaspaces XAP, Hazelcast, Apache Spark, Apache HBase and more



HARDWARE AND SOFTWARE REQUIREMENTS

Computer Requirements

- RAM: minimum 8 GB of RAM required for exercises and platform to operate, 16 GB and up recommended.
- Disk Space: At least 20 GB of free disk space
- Internet connection
- All machines connected to same Network

Supported Operating Systems

- Windows OS With VMware player running provided 123Completed-Ubuntu VM
- Some platforms are experienced on Windows and some on Ubuntu Linux (using the VM)

Additional Software Requirements

- PDF Reader
- Java JDK 7u55 (Install in a directory with a short path, without spaces)
- Zip software

Class HW required

- Projector 1024*768 minimum resolution
- White Board
- Erasable Markers
- Desktops or Laptops (see HW Requirements)
- 12-24 ports Switch
- Internet connectivity
- Electricity outlets for all computers/monitors and other equipment.
- At least 3 electricity outlets next to instructor location.

AGENDA	
Day 1: Lesson 1: Course Introduction	Duration: 30 minutes
<ul style="list-style-type: none"> • Course Introduction • Courseware walkthrough • Documentation • Lab 	
Day 1: Lesson 2: Introduction to Big Data and No SQL	Duration: 0.5 Day
<ul style="list-style-type: none"> • What is Big Data? • Big Data challenges and complexity • General concepts • Architecture considerations • Presenting use cases of internet companies (e.g. Facebook) • The Data Scientist • RDBMS: Advantages and disadvantages / Impedance Mismatch • No-SQL vs. Traditional Enterprise Relational Data: • CAP theorem vs. ACID / Dynamic schema, sharding, replications and caching / Performance • Scaling vs consistency • No-SQL types and use cases: Key/value stores / Document databases / Column oriented databases / Graph databases • When (not) to use No-SQL? 	
Day 1: Lesson 3: Introduction to Big Table Databases – Apache Cassandra DB	Duration: 0.5 Day
<ul style="list-style-type: none"> • Cassandra DB Introduction • Getting Started • Google Big Table • Amazon Dynamo • Cassandra Query Language Shell - CQLSH • Replication & Partitioning • Basic administration • Cluster configuration • Lab (Install Configure and Experience) 	
Day 2: Lesson 4: Quick Introduction to Document Databases – Mongo DB	Duration: 0.25 Day
<ul style="list-style-type: none"> • Mongo DB Introduction • Getting Started - Installation • Mongo DB basic commands • Aggregation • Replication • Sharding • Trainer Demo 	
Day 2: Lesson 5: Quick Introduction to Hadoop HBase DB	Duration: 0.25 Day
<ul style="list-style-type: none"> • HBase DB Introduction • Getting Started • HBase Data Model • HBase Shell and Basic Command • Physical Model • Architecture • Trainer Demo 	

Day 2: Lesson 6: Introduction to In Memory Data Grids - Redis	Duration: 0.5 Day
<ul style="list-style-type: none"> • What is Redis? • Why Redis? • Redis Data Store • Redis data structures • String • List • Hashset • Set • Sorted Set • Patterns • Lab 	
Day 3: Lesson 7: Introduction to Apache SPARK	Duration: 0.5 Day
<ul style="list-style-type: none"> • Apache SPARK Introduction • Getting Started • SPARK architecture • SPARK processing • Map Reduce • Example • Lab (Install Configure and Experience) 	
Day 3: Lesson 8: The Data Architect + Summary	Duration: 0.25 Day
<ul style="list-style-type: none"> • The Data Architect • Every Data element should be analyzed • Data Characteristics (Read rarely, Read once, Read many, Write Once, Write Many (Updated after inserted) • Read if Exist • Final Lab Session (Combine all the pieces) • Summary • Wrap Up 	
Day 3: Lesson 9: NoSQL DB Comparison	Duration: 0.25 Day
<ul style="list-style-type: none"> • NO SQL DBs General Comparison • Performance Comparison • When to use which? 	
Day 3: Lesson 10: Summary	Duration: 15 minutes
<ul style="list-style-type: none"> • Summary • Wrap Up 	