

INFOSOFT IT SOLUTIONS

Training | Projects | Placements

Revathi Apartments, Ameerpet, 1st Floor, Opposite Annapurna Block, Infosoft It solutions,
Software Training & Development Institute, +91 - 9059683947 | +91 - 9182540872

CASSANDRA TRAINING COURSE

Introduction to Big Data, and Cassandra

Topics:

- Introduction to Big Data and Problems caused by it
- 5V – Volume, Variety, Velocity, Veracity and Value
- Traditional Database Management System
- Limitations of RDMS
- NOSQL databases
- Common characteristics of NoSQL databases
- CAP theorem
- How Cassandra solves the Limitations?
- History of Cassandra
- Features of Cassandra

Hands On:

- Edureka VM tour

Cassandra Data Model

Topics:

- Introduction to Database Model
- Understand the analogy between RDBMS and Cassandra Data Model
- Understand following Database Elements: *Cluster, Keyspace, Column Family/Table, Column*
- Column Family Options
- Columns
- Wide Rows, Skinny Rows
- Static and dynamic tables

Hands-On:

- Creating Keyspace
- Creating Tables

Cassandra Architecture

Topics:

- Cassandra as a Distributed Database
- Key Cassandra Elements
 - a. Memtable*
 - b. Commit log*
 - c. SSTables*
- Replication Factor

- Data Replication in Cassandra
- Gossip protocol – Detecting failures
- Gossip: Uses
- Snitch: Uses
- Data Distribution
- Staged Event-Driven Architecture (SEDA)
- Managers and Services
- Virtual Nodes: Write path and Read path
- Consistency level
- Repair
- Incremental repair

Deep Dive into Cassandra Database

Topics:

- Replication Factor
- Replication Strategy
- Defining columns and data types
- Defining a partition key
- Recognizing a partition key
- Specifying a descending clustering order
- Updating data
- Tombstones
- Deleting data
- Using TTL
- Updating a TTL

Hands-on/Demo

- Create Keyspace in Cassandra
- Check Created Keyspace in System_Schema.Keyspaces
- Update Replication Factor of Previously Created Keyspace
- Drop Previously Created Keyspace
- Create A Table Using cqlsh
- Create A Table Using UUID & TIMEUUID
- Create A Table Using Collection & UDT Column
- Create Secondary Index On a Table
- Insert Data Into Table
- Insert Data into Table with UUID & TIMEUUID Columns
- Insert Data Using COPY Command
- Deleting Data from Table

Node Operations in a Cluster

Topics:

- Cassandra nodes
- Specifying seed nodes
- Bootstrapping a node
- Adding a node (Commissioning) in Cluster
- Removing (Decommissioning) a node
- Removing a dead node
- Repair
- Read Repair
- What's new in incremental repair

- Run a Repair Operation
- Cassandra and Spark Implementation

Hands On:

- Commissioning a Node
- Decommissioning a Node
- Nodetool Commands

Managing and Monitoring the Cluster

Topics:

- Cassandra monitoring tools
- Logging
- Tailing
- Using Nodetool Utility
- Using JConsole
- Learning about OpsCenter
- Runtime Analysis Tools

Hands On:

- JMX and Jconsole
- OpsCenter

Backup & Restore and Performance Tuning

Topics:

- Creating a Snapshot
- Restoring from a Snapshot
- RAM and CPU recommendations
- Hardware choices
- Selecting storage
- Types of Storage to Avoid
- Cluster connectivity, security and the factors that affect distributed system performance
- End-to-end performance tuning of Cassandra clusters against very large data sets
- Load balance and streams

Hands On:

- Creating Snapshots
- Integration with Kafka
- Integration with Spark

Hosting Cassandra Database on Cloud

Topics:

- Security
- Ongoing Support of Cassandra Operational Data
- Hosting a Cassandra Database on Cloud

Hands On:

- Hosting Cassandra Database on Amazon Web Services