

Cassandra Comprehensive Training

Pre-requisites:

- Knowledge on Core Java
- Basic understanding on Linux OS (Ubuntu)

Projects:

Two mini-projects (Web Based) in Java will be performed in this during

based on coverage to give a participant a better insight on Cassandra.

Daywise Agenda:

Introduction to Cassandra

- Introducing Cassandra
- Understanding what cassandra is?
- Learning what cassandra is used for?
- CAP Theorem
- Cluster Architecture
- Eventual Consistency
- Understanding System Requirements
- Understanding our lab

Getting Started with Cassandra

- Understanding Cassandra as Distributed DB
- Snitch
- Gossip
- Learning How Data gets distributed
- Replication
- Virtual Nodes

Installing Cassandra

- Downloading Cassandra
- Java
- Understanding cassandra configuration files
- Cassandra foreground and background mode
- Checking Cassandra Status
- Accessing and understanding Log Structure

Communicating with Cassandra

- Using CQLSH
- Creating a Database
- Defining a Keyspace
- Deleting a Keyspace
- Creating a Table
- Defining Columns and Datatypes
- Defining Primary Key

- Recognizing a Partition Key

- Specifying a descending cluster order
- Understanding ways to write data
- Using INSERT INTO command
- Using COPY command
- Understanding how data is stored in Cassandra
- Understanding How data is stored in Disk

Understanding Data Modelling in Cassandra

- Understanding Data model
- Understanding where clause criteria in Cassandra
- Loading Bulk Data
- JSON format Import and Export
- Using Primary Index
- Creating a Secondary Index
- Defining a Composite Partition Key

Creating an Application using Cassandra Backend

- Understanding Cassandra Drivers
- Exploring the Datastax Java Driver
- Setting up Eclipse Environment
- Creating an Application WebPage
- Acquiring Java Driver Files
- Understanding Packaging using Maven
- Understanding Packaging using Manual Methods
- Connecting to Cassandra Cluster using WebPage
- Executing a Query using WebPage in Cassandra
- Using MVC Pattern Example

Updating and Deleting Data

- Updating Data
- Understanding How updating Works
- Deleting Data
- Understanding the role of Tombstones
- Using TTL

Cassandra Multinode Cluster Setup

- Understanding Hardware Choices for production
- Understanding RAM and CPU Recommendations
- Things to be considered while Selecting storage
- Things to be considered while Deploying in Cloud
- Understanding Cassandra Nodes
- Network Connection Setup
- Specifying Seed Nodes
- Bootstrapping a node
- Cleaning up a node

- Using cassandra-stress for stress testing cluster

Cassandra Monitoring and Maintenance --- PART 1

- Understanding Cassandra Monitoring Tools
- Using Nodetool
- Using Jconsole
- Learning about OpsCenter
- Understanding Repair
- Repairing Nodes
- Understanding Consistency
- Understanding Hinted Handoff
- Understanding Read Repair

Cassandra Monitoring and Maintenance --- PART 2

- Removing a node
- Putting a node back to service
- Decommissioning a node
- Removing a dead node
- Redefining Multiple Data centers
- Changing Snitch Types
- Modifying cassandra-rackdc.properties
- Changing Replication Strategy

Understanding Backup, Restore and Performance Tuning

- Understanding Backup & Restore Concepts in Cassandra
- Taking a Snapshot
- Incremental Backup
- Using Commit Log Feature
- Using Restore Methods
- Storage Strategies and OS tuning
- JVM Tuning
- Caching Strategies
- Compaction and Compression
- Stress Testing Strategies

System Requirements:

- Intel Corei3 Processor or later
- 8GB RAM
- 100GB free HDD space
- Windows 7 or later OS