

Cassandra – Developer Track

Duration

2day

Introducing Cassandra

- A Quick Review of Relational Databases
- Use Cases for Cassandra
 - Large Deployments
 - Lots of Writes, Statistics, and Analysis
 - Geographical Distribution
 - Evolving Applications

Installing Cassandra

- Basic CLI Commands
- Connecting to a Server
- Describing the Environment
- Creating a Keyspace and Column Family
- Writing and Reading Data

The Cassandra Data Model

- The Relational Data Model
- Clusters
- Keyspaces
- Column Families
 - Column Family Options
- Columns
 - Wide Rows, Skinny Rows
 - Column Sorting
- Super Columns
 - Composite Keys
- Design Differences Between RDBMS and Cassandra
 - No Query Language
 - No Referential Integrity
 - Secondary Indexes
 - Sorting Is a Design Decision
 - Denormalization
- Design Patterns
 - Materialized View
 - Valueless Column
 - Aggregate Key

The Cassandra Architecture

- System Keyspace
- Peer-to-Peer
- Gossip and Failure Detection
- Anti-Entropy and Read Repair
- Memtables, SSTables, and Commit Logs
- Compaction
- Bloom Filters

Configuring Cassandra

- Keyspaces
 - Defining Keyspaces
- Creating a Column Family
 - Designing Column Families
- Replicas
- Replica Placement Strategies
 - Indexes in Cassandra
 - Primary
 - Secondary

Planning Your Data Model

Managing and Accessing Data in Cassandra

- About Writes in Cassandra
- Managing Stored Data
- About Transactions and Concurrency Control
- About Inserts and Updates
- About Deletes
- Simple Get
- Batches Deletes

Cassandra CLI

Cassandra web console

Java: Hector Client API

- Initializing a cluster
- Creating a template
- Accessing data
- Hector Pool Architecture
- Using CQL

Introduction to Pelops (Java)

Monitoring

- Logging
- Cassandra MBeans
- Runtime Analysis Tools

Maintenance

- Ring Information
- Snapshot