

Introduction to Cassandra

Code:	TTDS6776
Length:	3 days
URL:	View Online

The Cassandra (C*) database is a massively scalable NoSQL database that provides high availability and fault tolerance, as well as linear scalability when adding new nodes to a cluster. It has many powerful capabilities, such as tunable and eventual consistency, that allow it to meet the needs of modern applications, but also introduce a new paradigm for data modeling that many organizations do not have the expertise to use in the best way.

Introduction to Cassandra is a hands-on course designed to teach attendees the basics of how to create good data models with Cassandra. This technical course has a focus on the practical aspects of working with C*, and introduces essential concepts needed to understand Cassandra, including enough coverage of internal architecture to make good decisions. It is hands-on, with labs that provide experience in core functionality. Students will also explore CQL (Cassandra Query Language), as well as some of the “anti-patterns” that lead to non-optimal C* data models and be ready to work on production systems involving Cassandra.

Skills Gained

- Understand the Big Data needs that C* addresses
- Be familiar with the operation and structure of C*
- Be able to install and set up a C* database
- Use the C* tools, including cqlsh, nodetool, and ccm (Cassandra Cluster Manager)
- Be familiar with the C* architecture, and how a C* cluster is structured
- Understand how data is distributed and replicated in a C* cluster
- Understand core C* data modeling concepts, and use them to create well-structured data models
- Be familiar with the C* eventual consistency model and use it intelligently
- Be familiar with consistency mechanisms such as read repair and hinted handoff
- Understand and use CQL to create tables and query for data
- Know and use the CQL data types (numerical, textual, uuid, etc.)
- Be familiar with the various kinds of primary keys available (simple, compound, and composite primary keys)
- Be familiar with the C* write and read paths
- Understand C* deletion and compaction

Prerequisites

Attendees should have incoming experience with and knowledge of SQL. Some familiarity with distributed systems is also helpful.

Course Details

Course Outline

- Session 1: Cassandra Overview
 - Session 2: Cassandra Architecture and CQL Overview
 - Session 3: Data Modeling and CQL Core Concepts
 - Session 4: Additional CQL Capabilities
 - Session 5: Data Consistency In Cassandra
 - Session 6: Internal Mechanisms
 - Session 7: Working with IntelliJ
-

Download Whitepaper: Accelerate Your Modernization Efforts with a Cloud-Native Strategy

Get Your Free Copy Now