

Introduction to Cassandra (TTDS6776)

Category: AI, Data & Analytics | **Vendor:** Trivera Tech

Duration: 24.00 hours (3 days)

19.5 CPD Hours

Rating: ★ 4.6 (5,878 reviews)

Course Information

Delivery Format: Instructor Led - Online

Course Overview

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.

About This Course

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.

Who Should Attend

This course is deal for software developers, database administrators, and data architects interested in learning about distributed database management systems for large-scale data processing.

Prerequisites & Entry Requirements

General Prerequisites:

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

Learning Outcomes

Upon successful completion of this course, participants will be able to:

The goal of this course is to enable technical students new to Cassandra to begin working with Cassandra in an optimal manner. Throughout the course students will learn to:

- 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

Detailed Course Outline

Session 1: Cassandra Overview

- Why We Need Cassandra - Big Data Challenges vs RDBMS
- High level Cassandra Overview
- Cassandra Features
- Optional: Basic Cassandra Installation and Configuration

Session 2: Cassandra Architecture and CQL Overview

- Cassandra Architecture Overview
- Cassandra Clusters and Rings
- Nodes and Virtual Nodes
- Data Replication in Cassandra
- Introduction to CQL
- Defining Tables with a Single Primary Key
- Using cqlsh for Interactive Querying
- Selecting and Inserting/Upserting Data with CQL
- Data Replication and Distribution
- Basic Data Types (including uuid, timeuuid)

Session 3: Data Modeling and CQL Core Concepts

- Defining a Compound Primary Key
- CQL for Compound Primary Keys
- Partition Keys and Data Distribution
- Clustering Columns
- Overview of Internal Data Organization
- Overview of Other Querying Capabilities
- ORDER BY, CLUSTERING ORDER BY, UPDATE, DELETE, ALLOW FILTERING
- Batch Queries
- Data Modeling Guidelines
- Denormalization
- Data Modeling Workflow
- Data Modeling Principles
- Primary Key Considerations
- Composite Partition Keys
- Defining with CQL
- Data Distribution with Composite Partition Key
- Overview of Internal Data Organization

Session 4: Additional CQL Capabilities

- Indexing
- Primary/Partition Keys and Pagination with token()
- Secondary Indexes and Usage Guidelines
- Cassandra collections
- Collection Structure and Uses
- Defining and Querying Collections (set, list, and map)
- Materialized View
- Overview
- Usage Guidelines

Session 5: Data Consistency In Cassandra

- Overview of Consistency in Cassandra
- CAP Theorem
- Eventual (Tunable) Consistency in C* - ONE, QUORUM, ALL
- Choosing CL ONE
- Choosing CL QUORUM
- Achieving Immediate Consistency
- Overview of Other Consistency Levels
- Supportive Consistency Mechanisms
- Writing / Hinted Handoff
- Read Repair
- Nodetool repair

Session 6: Internal Mechanisms

- Ring Details
- Partitioners
- Gossip Protocol
- Snitches
- Write Path
- Overview / Commit Log
- Memtables and SSTables
- Write Failure
- Unavailable Nodes and Node Failure
- Requirements for Write Operations
- Read Path Overview
- Read Mechanism
- Replication and Caching
- Deletion/Compaction Overview
- Delete Mechanism
- Tombstones and Compaction

Session 7: Working with IntelliJ

- Configuring JDBC Data Source for Cassandra
- Reading Schema Information
- Querying and Editing Tables.

, Session 1: Cassandra Overview

Why We Need Cassandra - Big Data Challenges vs RDBMS

High level Cassandra Overview

Cassandra Features

Optional: Basic Cassandra Installation and Configuration

Session 2: Cassandra Architecture and CQL Overview

Cassandra Architecture Overview

Cassandra Clusters and Rings

Nodes and Virtual Nodes

Data Replication in Cassandra

Introduction to CQL

Defining Tables with a Single Primary Key

Using cqlsh for Interactive Querying

Selecting and Inserting/Upserting Data with CQL

Data Replication and Distribution

Basic Data Types (including uuid, timeuuid)

Session 3: Data Modeling and CQL Core Concepts

Defining a Compound Primary Key

CQL for Compound Primary Keys

Partition Keys and Data Distribution

Clustering Columns

Overview of Internal Data Organization

Overview of Other Querying Capabilities

ORDER BY, ALLOW FILTERING

Batch Queries

Data Modeling Guidelines

Denormalization

Data Modeling Workflow

Data Modeling Principles

Primary Key Considerations

Composite Partition Keys

Defining with CQL

Data Distribution with Composite Partition Key

Overview of Internal Data Organization

Session 4: Additional CQL Capabilities

Indexing

Primary/Partition Keys and Pagination with token()

Secondary Indexes and Usage Guidelines

Cassandra collections

Collection Structure and Uses

Defining and Querying Collections (set, and map)

Materialized View

Overview

Usage Guidelines

Session 5: Data Consistency In Cassandra

Overview of Consistency in Cassandra

CAP Theorem

Eventual (Tunable) Consistency in C* - ONE, ALL

Choosing CL ONE

Choosing CL QUORUM

Achieving Immediate Consistency

Overview of Other Consistency Levels

Supportive Consistency Mechanisms

Writing / Hinted Handoff

Read Repair

Nodetool repair

Session 6: Internal Mechanisms

Ring Details

Partitioners

Gossip Protocol

Snitches

Write Path

Overview / Commit Log

Memtables and SSTables

Write Failure

Unavailable Nodes and Node Failure

Requirements for Write Operations

Read Path Overview

Read Mechanism

Replication and Caching

Deletion/Compaction Overview

Delete Mechanism

Tombstones and Compaction

Session 7: Working with IntelliJ

Configuring JDBC Data Source for Cassandra

Reading Schema Information

Querying and Editing Tables.

Additional Course Details

Nexus Humans Introduction to Cassandra (TTDS6776) training program is a workshop that presents an invigorating mix of sessions, lessons, and masterclasses meticulously crafted to propel your learning expedition forward.

This immersive bootcamp-style experience boasts interactive lectures, hands-on labs, and collaborative hackathons, all strategically designed to fortify fundamental concepts.

Guided by seasoned coaches, each session offers priceless insights and practical skills crucial for honing your expertise. Whether you're stepping into the realm of professional skills or a seasoned professional, this comprehensive course ensures you're equipped with the knowledge and prowess necessary for success.

While we feel this is the best course for the Introduction to Cassandra (TTDS6776) course and one of our Top 10 we encourage you to read the course outline to make sure it is the right content for you. Additionally, private sessions, closed classes or dedicated events are available both live online and at our training centres in Dublin and London, as well as at your offices anywhere in the UK, Ireland or across EMEA.

Frequently Asked Questions

Q: What delivery options are available for Introduction to Cassandra (TTDS6776)?

We offer multiple delivery formats:

- Live Instructor-Led Classroom Online (Virtual/Live Online)
 - Traditional Instructor-Led Classroom Training (ILT)
 - On-site delivery at your offices anywhere in United Kingdom
 - Private dedicated courses customized for your team
-

Q: How many CPD hours does this course provide?

The 3-day Introduction to Cassandra (TTDS6776) course provides up to 19.5 CPD hours of structured learning. CPD certificates can be provided upon request.

Q: What is the duration of the Introduction to Cassandra (TTDS6776) training?

The training takes place over 3 day(s), with each day lasting approximately 24.00 hours including breaks for lunch and refreshments.

Q: Do you provide corporate training for Introduction to Cassandra (TTDS6776)?

Yes, we provide corporate training, dedicated training, and closed classes for Introduction to Cassandra (TTDS6776). Training can take place anywhere in United Kingdom including London, Manchester, Birmingham, Edinburgh, or live online allowing teams from across United Kingdom or internationally to attend.

Q: What related terms do people search for?

Popular related searches include: Cassandra; Databases; Data Science

Q: Why choose Nexus Human for Introduction to Cassandra (TTDS6776)?

Nexus Human is recognized as one of the leading training providers. Our trainers have won multiple awards including:

- Small Firms Best Trainer Award
- National Training Partner of the Year (Ireland) - Multiple Years
- Global Top 30 Instructor Awards (2012, 2019, 2021)
- Tech Excellence Award Nominations
- Learning Performance Institute (LPI) External Training Provider Sponsor 2024

Q: Are there any discount codes available?

Yes! Use discount code **PENPAL5** when booking your Introduction to Cassandra (TTDS6776) training. Please note that only one discount code can be used per booking and cannot be combined with other special offers.

Nexus Human

Professional Training & Development

 Email: info@nexushuman.com

 Website: www.nexushuman.com

 Phone: +353 1 XXX XXXX (Ireland) | +44 20 XXXX XXXX (UK)