

From 0 to 1: The Cassandra Distributed Database

Modality: Self-Paced Learning

Duration: 6 Hours

SATV Value:

CLC:

NATU:

SUBSCRIPTION: Learn, Master

About this course:

Has your data gotten huge, unwieldy and hard to manage with a traditional database? Is your data unstructured with an expanding list of attributes? Do you want to ensure your data is always available even with server crashes? Look beyond Hadoop - the Cassandra distributed database is the solution to your problems. This course helps you set up a cluster with multiple nodes to distribute data across machines. Cassandra is a columnar store. There are no empty cells or space wasted when you store data with variable and expanding attributes. Cassandra uses partitioning and replication to ensure that your data is available even when nodes in a cluster go down.

The average salary for Database Engineer is **\$86,509** per year.

Course Objective:

After completing this course, students will have a working understanding of:

- Introduction: Cassandra as a distributed, decentralized, columnar store
- The Cassandra Cluster Manager
- The Cassandra Data Model
- Shell Commands
- Keys And Indexes: Primary Keys, Partition Keys, Clustering Key, Secondary Indexes
- Storage Systems

Audience:

This course is intended for:

- Engineers and analysts who understand traditional, relational databases and want to move to big data storage systems
- Students who are just starting out understanding databases and have no prior experience with one

Prerequisites:

- The basics of SQL and traditional relational databases

- The basics of Java in order to use the Cassandra Java library

Suggested prerequisites courses:

- [SQL Database for Beginners](#)
- [Learn to Program in Java](#)

Course Outline:

- You, This Course and Us
- Introduction: Cassandra as a distributed, decentralized, columnar store
- Install And Set Up
- The Cassandra Cluster Manager
- The Cassandra Data Model
- Shell Commands
- Keys And Indexes: Primary Keys, Partition Keys, Clustering Key, Secondary Indexe
- Tunable Consistency
- Storage Systems
- A Mini-Project: A Miniature Catalog Management System In Java