


DS320: DataStax Enterprise Analytics with Apache Spark™

Instructor-Led Training

Overview

In this course, you will learn how to effectively and efficiently solve analytical problems with DataStax Enterprise Analytics. You will learn about Apache Cassandra™ and Apache Spark, including Spark API, Spark-Cassandra Connector, Spark SQL, Spark Streaming, and crucial performance optimization techniques.



Duration	3 days
Audience	Software professionals proficient with Cassandra core concepts and looking to learn DataStax Enterprise Analytics
Prerequisites	Completion of the DataStax Enterprise Foundations of Apache Cassandra™ online course. Linux command line and programming experience.

Course details

- Introduction
 - Data Analytics with Cassandra and Spark
 - Spark Architecture
 - Spark Shell
 - Spark Web UI
- Spark Essentials
 - A Spark "Hello, World!"
 - Resilient Distributed Dataset
 - Creating an RDD
 - RDD Transformations
 - RDD Actions
- Connecting Spark to Cassandra
- Tuning and Partitioning
 - Understanding Partitioning in Spark
 - Partitioning Rules
 - Controlling Partitioning
 - Data Shuffling
- Optimizing with the Spark/Cassandra Connector
 - Counting Rows in a Cassandra Table
 - Grouping Rows Using Cassandra
 - Joining Cassandra Tables
 - Cassandra-Aware Partitioning
- Spark Streaming
 - Discretized Stream

About DataStax

It starts with a human desire, and when a universe of technology, devices and data aligns, it ends in a moment of fulfillment and insight. Billions of these moments occur each second around the globe. They are moments that can define an era, launch an innovation, and forever alter for the better how we relate to our environment. DataStax is the power behind the moment. Built on the unique architecture of Apache Cassandra™, DataStax Enterprise is the always-on data platform and has been battle-tested for the world's most innovative, global applications.

With more than 500 customers in over 50 countries, DataStax provide data management to the world's most innovative companies, such as Netflix, Safeway, ING, Adobe, Intuit, Target and eBay. Based in Santa Clara, Calif., DataStax is backed by industry-leading investors including Comcast Ventures, Crosslink Capital, Lightspeed Venture Partners, Kleiner Perkins Caufield & Byers, Meritech Capital, Premji Invest and Scale Venture Partners. For more information, visit DataStax.com/customers or follow us on @DataStax.

- Retrieving Data from Cassandra
- Working with Cassandra Rows
- Converting Cassandra Rows to Scala Tuples and Objects
- Saving Data to Cassandra
- Optimization is not Always Premature
 - Broadcast Variables
 - Accumulator Variables
 - RDD Persistence
- The Utility of Key/Value Pairs
 - Key-Value Pair RDD
 - Key-Value Pair RDDs: Aggregation
 - Key-Value Pair RDDs: Grouping and Sorting
 - Key-Value Pair RDDs: Joins
 - Key-Value Pair RDDs: Set Operations
- Spark Streaming Architecture
 - Your First Spark Streaming Application
 - DStream Stateless Transformations
 - DStream Stateful Transformations
 - DStream Window Transformations
 - DStream Output Operations
 - Checkpointing and Recovery
 - DStream Persistence
 - Controlling Parallelism
- Spark SQL
 - Spark SQL Basics
 - Creating DataFrames
 - Accessing DataFrame Schema and Rows
 - RDD Operations on DataFrames
 - Language-Integrated Queries
 - Saving DataFrames to Cassandra
 - Querying Cassandra with SQL
 - Writing Efficient SQL Queries

Benefits

Classroom training is delivered by DataStax experts, who will share first-hand experiences and practical use cases. A high level of interactivity with the trainer coupled with hands-on labs deliver an active learning experience.

A private, on-site training environment provides the option to tailor the delivery based on your needs and preferences. We will work with your team to determine areas requiring emphasis during delivery. Additionally, on-site training fosters teamwork and team learning that continue through the project.

For more information, please contact your account executive or send an email to training@datastax.com.

Terms

Please refer to the DataStax Services Terms for specific details regarding terms of Training delivery.

- This is a pre-paid Service offering.
- The Service expires 365 days after receipt of a valid customer order.
- The maximum class size is 20 attendees.
- This training is delivered onsite.
- A virtual classroom option is available, but we do not recommend mixing in-person and remote attendees. The remote attendees have limited learning experience in this environment.

About DataStax

It starts with a human desire, and when a universe of technology, devices and data aligns, it ends in a moment of fulfillment and insight. Billions of these moments occur each second around the globe. They are moments that can define an era, launch an innovation, and forever alter for the better how we relate to our environment. DataStax is the power behind the moment. Built on the unique architecture of Apache Cassandra™, DataStax Enterprise is the always-on data platform and has been battle-tested for the world's most innovative, global applications.

With more than 500 customers in over 50 countries, DataStax provide data management to the world's most innovative companies, such as Netflix, Safeway, ING, Adobe, Intuit, Target and eBay. Based in Santa Clara, Calif., DataStax is backed by industry-leading investors including Comcast Ventures, Crosslink Capital, Lightspeed Venture Partners, Kleiner Perkins Caufield & Byers, Meritech Capital, Premji Invest and Scale Venture Partners. For more information, visit DataStax.com/customers or follow us on @DataStax.