



“Cloudera’s instructor was excellent, offering clear and concise training that was easy to understand. His wide-ranging peripheral knowledge helped apply the course materials to real-world situations. I look forward to attending another course.”

Comscore

Cloudera Training for Apache HBase

Take your knowledge to the next level with Cloudera’s Apache Hadoop Training and Certification

Cloudera University’s three-day training course for Apache HBase enables participants to store and access massive quantities of multi-structured data and perform hundreds of thousands of operations per second.

Advance Your Ecosystem Expertise

Apache HBase is a distributed, scalable, NoSQL database built on Apache Hadoop. HBase can store data in massive tables consisting of billions of rows and millions of columns, serve data to many users and applications in real time, and provide fast, random read/write access to users and applications.

Hands-On Hadoop

Through instructor-led discussion and interactive, hands-on exercises, participants will navigate the Hadoop ecosystem, learning topics such as:

- The use cases and usage occasions for HBase, Hadoop, and RDBMS
- Using the HBase shell to directly manipulate HBase tables
- Designing optimal HBase schemas for efficient data storage and recovery
- How to connect to HBase using the Java API to insert and retrieve data in real time
- Best practices for identifying and resolving performance bottlenecks

Audience & Prerequisites

This course is appropriate for developers and administrators who intend to use HBase. Prior experience with databases and data modeling is helpful, but not required. Knowledge of Java is assumed. Prior knowledge of Hadoop is not required, but Cloudera Developer Training for Apache Hadoop provides an excellent foundation for this course.

HBase Certification

Upon completion of the course, attendees are encouraged to continue their study and register for the Cloudera Certified Specialist in Apache HBase (CCSHB) exam. Certification is a great differentiator; it helps establish you as a leader in the field, providing employers and customers with tangible evidence of your expertise.

Course Outline: Cloudera Training for Apache HBase

Introduction

Introduction to Hadoop and HBase

- What Is Big Data?
- Introducing Hadoop
- Hadoop Components
- What Is HBase?
- Why Use HBase?
- Strengths of HBase
- HBase in Production
- Weaknesses of HBase

HBase Tables

- HBase Concepts
- HBase Table Fundamentals
- Thinking About Table Design

The HBase Shell

- Creating Tables with the HBase Shell
- Working with Tables
- Working with Table Data

HBase Architecture Fundamentals

- HBase Regions
- HBase Cluster Architecture
- HBase and HDFS Data Locality

HBase Schema Design

- General Design Considerations
- Application-Centric Design
- Designing HBase Row Keys
- Other HBase Table Features

Basic Data Access with the HBase API

- Options to Access HBase Data
- Creating and Deleting HBase Tables
- Retrieving Data with Get
- Retrieving Data with Scan
- Inserting and Updating Data
- Deleting Data

More Advanced HBase API Features

- Filtering Scans
- Best Practices
- HBase Coprocessors

HBase on the Cluster

- How HBase Uses HDFS
- Compactions and Splits

HBase Reads and Writes

- How HBase Writes Data
- How HBase Reads Data
- Block Caches for Reading

HBase Performance Tuning

- Column Family Considerations
- Schema Design Considerations
- Configuring for Caching
- Dealing with Time Series and Sequential Data
- Pre-Splitting Regions

HBase Administration and Cluster Management

- HBase Daemons
- ZooKeeper Considerations
- HBase High Availability
- Using the HBase Balancer
- Fixing Tables with hbck
- HBase Security

HBase Replication and Backup

- HBase Replication
- HBase Backup
- MapReduce and HBase Clusters

Using Hive and Impala with HBase

- Using Hive and Impala with HBase

Conclusion

Appendix A: Accessing Data with Python and Thrift

- Thrift Usage
- Working with Tables
- Getting and Putting Data
- Scanning Data
- Deleting Data
- Counters
- Filters

Appendix B: OpenTSDB

Cloudera Certified Specialist in Apache HBase (CCSHB)

Establish yourself as a trusted and valuable resource by completing the certification exam for Apache HBase. CCSHB certifies your ability to design optimal HBase schemas and configure an HBase cluster for efficient data storage and recovery. The exam can be demanding and will test your fluency with concepts and terminology in the following areas:

Core HBase Concepts

Recognize the fundamental characteristics of HBase and its role in a big data ecosystem. Identify the differences between Apache HBase and a traditional RDBMS. Describe the relationship between HBase and HDFS.

Data Model

Describe how an HBase table is physically stored on disk. Identify the differences between a Column Family and a Column Qualifier. Describe how HBase cells store data. Detail what happens to data when it is deleted

Architecture

Identify the major components of an HBase cluster. Recognize how regions work and their benefits under various scenarios. Describe how a client finds a row in an HBase table. Understand the function and purpose of minor and major compactions. Describe Region-Server splits

Schema Design

Describe the factors to be considered with creating Column Families. Define the row keys for optimal read performance and locality in an access pattern

API

Describe the functions and purpose of the HBaseAdmin class. Given a table and rowkey, use the get() operation to return specific versions of that row. Describe the behavior of the checkAndPut() method

Administration

Recognize how to create, describe and access data in tables from the shell. Describe how to bulk load data into HBase. Recognize the benefits of managed region splits