COD-283

Developing Applications with Cloudera Operational Database

Course Overview

Course Type
Instructor-led training

Level
Intermediate

Duration
3 days

Platform
Cloudera

Topics Covered

- Apache HBase
- Apache Phoenix
- Apache OMID

About This Training

Take your knowledge to the next level with "Developing Applications with Cloudera Operational Database" three-day training course enables participants to store and access massive quantities of multi-structured data and perform hundreds of thousands of operations per second.

Business Values

- **Operational Speed at Scale**
  Cloudera Operational Database enables applications to store and access large amounts of data and allow processing hundreds of thousands of operations per second.
  For customers, this translates into faster application response times and improving the experience for users.

- **Reduced Deployment Time**
  Taking this training enables customers to build efficient schema layouts and application design faster. The curriculum displays best practices and common mistakes and bottlenecks.
  This allows customers to create faster and better applications.

- **Integration with Cloudera Data Platform**
  This course embeds Cloudera Operational Database into the larger Cloudera Data Platform. The course shows integration with Data Engineering and other Cloudera components.
  For customers, this integration removes the expensive and time-consuming needs to integrate Cloudera Operational Database with third-party or external frameworks, as Cloudera Data Platform already provides such components.

COD-283

Developing Applications with Cloudera Operational Database

Industry Use Cases

- **Operational Speed at Scale**
  Cloudera Operational Database allows for handling millions of financial transactions per second. This enables real-time applications such as fraud detection, monitoring and customer 360 profiles.

- **Telecommunications**
  Telcos generate large amounts of Call Detail Records (CDRs), network events, and subscriber data. Cloudera Operational Database can store and query such data for billing systems, network analytics and churn prediction.

- **Retail and E-Commerce**
  Real-time personalization, click-stream analytics and more low-latency applications can be build using Cloudera Operational Database.

- **Energy and Utilities**
  Smart grids and IoT sensor networks generate large amounts of continuous streams of time-series data. Cloudera Operational Database excels at ingesting, storing and querying this data for predictive maintenance, usage analytics and outage detection.

Hands-on Exercises

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
- How to use Phoenix to add SQL features and transactions on top of HBase

COD-283

Developing Applications with Cloudera Operational Database

Audience and prerequisites

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

This course is mostly appropriate for developers who intend to use HBase and/or Phoenix. Prior experience with databases and data modeling is helpful, but not required. Knowledge of Java is helpful but not required. Knowledge of Unix command line tools is helpful but not required. Prior knowledge of Hadoop is not required, but Cloudera Developer Training for Spark and Hadoop provides an excellent foundation for this course.

Also, administrators will find this course useful. The content will explain the underlying architecture. This will allow administrators to size the clusters correctly and monitor HBase tables and indices.

Students need to be able to establish outgoing ssh connections to perform the exercises of this course.

Advance your ecosystem expertise

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

Apache Phoenix will add SQL features and transactions on top of Apache HBase to allow for convenient data access and management.

COD-283

Developing Applications with Cloudera Operational Database

Introduction

- Team activity: Team Introductions

Operational Database

- Overview
- Others
- Use Cases
- Industries
- Architecture

HBase Shell

- HBase Shell Basics
- Running the HBase Shell
- Basic Commands
- Shell Command Syntax
- Shell Scripting
- Exercise: HBase Shell

HBase Import

- HBase Import Flow
- Exercise: HBase Import

HBase Logical Table Model

- Fundamental HBase Concepts
- Table Properties
- Example Application Data
- Rows
- Columns
- Columns and Column Families
- Conceptual View
- Storing Data in Tables
- Table Storage On Disk: Basics
- Data Storage Within a Column Family
- Read Operations
- Modify Operations
- Increment Operations
- Row Key
- RDBMS vs. HBase
- Replacing RDBMs with HBase
- Comparison with RDBMS Table Design
- HBase vs. RDBMS Schema Design
- Exercise: HBase Logical Table Model

© 2026 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies.
Information is subject to change without notice.

COD-283

Developing Applications with Cloudera Operational Database

HBase Data Access

- Gets and Puts
- Inserting Rows
- Retrieving Rows
- Scans
- Scanning
- Scanning Rows
- Removing Rows
- Deleting Data
- What Really Happens When You Alter Data
- Versions
- Exercise: HBase Data Access

HBase Write Path

- Durable Storage
- Store Files
- Client Put
- WAL
- MemStore
- Write Cache
- Flush
- Options
- HLog Files
- HFiles
- Recovery
- Minor Compaction
- Recommendations
- Splits
- Exercise: HBase Write Path

HBase Read Path

- Client Get
- Reading Data
- Block Cache
- Caching
- Caching Operations
- Configuration
- Bucket Cache
- Disabling Cache
- Caching Priorities

HBase Optimization

- Column Families Design
- Column Families Attributes
- Compression
- Block size
- Bloom Filters

© 2026 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies.
Information is subject to change without notice.

COD-283

Developing Applications with Cloudera Operational Database

HBase Java API

- Accessing Data in HBase
- HBase Java API
- HBase and Byte Arrays
- Simple Examples Using the HBase Java API (1)
- Simple Examples Using the HBase Java API (2)
- Using the Administration API to Create a Table
- Using the Administration API to Delete a Table
- Connecting to HBase
- Retrieving Rows
- Best Practice: Avoid Overuse of toBytes()
- Getting Previous Versions of a Cell in HBase
- Scans
- Reducing Scan Results Returned
- Scanner Caching
- Adding and Updating Data
- Deleting Data
- Exercise: HBase Java API

Phoenix Overview

- Features
- Resources
- Diagram
- Architecture
- Deployment

Phoenix Basics

- Command Line Client
- Metadata SQL Commands
- Creating a table
- Table Architecture
- Modifying and Deleting rows
- Reading data
- Exercise: Basic Syntax

Phoenix Transactions

- Overview
- Deployment
- Creating Tables
- Transaction Control
- Resources
- Exercise: Transactions

Phoenix Transactions

- Overview
- Deployment
- Creating Tables
- Transaction Control
- Resources
- Exercise: Transactions

Phoenix Bulk Loading

- Requirements
- Importing
- Exercise: Bulk Loading

© 2026 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies.
Information is subject to change without notice.

COD-283

Developing Applications with Cloudera Operational Database

Phoenix Secondary Indexes

- Overview
- Diagram
- Exercise: Secondary Indexes

Phoenix Views

- Overview
- Syntax
- Updatable Views
- Indexed Views
- Exercise: Views

Phoenix Salted Tables

- Overview
- Syntax
- Exercise: Salted Tables

© 2026 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies.
Information is subject to change without notice.

VERSION:
2.0.0