

ADMIN-237

Apache Ozone on Cloudera Base on Premises



Course Overview

Course Type

Instructor-led training

Level

Intermediate

Duration

4 days

Platform

Cloudera Base on Premises

Topics Covered

- Ozone Concepts
- Installation and Configuration
- Security
- File System Best Practices
- Maintenance
- Monitoring
- Performance Tuning
- Application Integration
- Data Migration from HDFS

About This Training

This all-round course by Cloudera Education introduces Apache Ozone, the next-generation hybrid storage service offering versatility and out-of-the-box compatibility without the limitations of HDFS. As the main developers of Apache Ozone, Cloudera have advised on this courseware to cover all aspects of Apache Ozone. This course covers a wide range of practical topics, such as Ozone architecture, internal operations, installation, file system usage, best practices, security, maintenance, monitoring, tuning, testing, application integration, and data migration.

What Skills You Will Gain

- Understanding the Benefits of Using Ozone
- Installing and Configuring Ozone with Security
- Managing Files and Objects in Ozone
- Performance Tuning and Doing Baseline Tests
- Controlling Replication and Understanding Failover and Recovery
- Performing Maintenance Tasks
- Monitoring Ozone Using Recon Service
- Integrating Hive, Impala, Spark, NiFi, and Flink with Ozone
- Migrating Data from HDFS to Ozone

Who Should Take This Course

This course is for new and existing customers using Cloudera Base on Premises who are interested in migrating their data and applications to Apache Ozone. The course is designed for Platform Administrators, Developers, Business Users, Data Engineers, and Business Operations teams. A general knowledge of Cloudera is beneficial but not required.

Other Training that Might Interest You

- [INTRO-001: Building Secure Clusters](#)
- [DOPS-241 Cloudera DataFlow: Flow Management with Apache NiFi](#)
- [DENG-254 Preparing with Cloudera Data Engineering and Apache Spark](#)
- [Learning Path: Data Analyst](#)

ADMIN-237

Apache Ozone on Cloudera Base on Premises



Ozone Introduction, Concepts, and Installation

- Introduction
- Why Ozone?
- Ozone Concepts
- Installation and Configuration
- Ozone Client Configuration

Security

- Security Overview
- Security Installation with Kerberos
- Securing Data using Apache Ranger
- Verifying Data Nodes are Secure

Working with Ozone

- File System and Admin Commands
- Transferring Data between Ozone and HDFS
- Setting Quotas and ACLs
- Controlling Replication
- DataNode Layout and Failover

Ozone Administration

- Maintenance Procedures
- Monitoring Ozone with Recon
- Integrating Ozone with Prometheus and Grafana

Performance Tuning

- Ozone Pipelines
- Baseline Load Testing with Freon
- Upgrading Ozone

Application Integration and Data Migration

- Ozone Data Locality
- Integrating Ozone with Apache Hive and Impala
- Integrating Ozone with Apache Spark
- Integrating Ozone with Apache NiFi

Data Migration

- HDFS and Ozone
- Ozone Configuration
- Mapping the HDFS Namespace to Ozone
- Data Migration using DistCp