

ADMIN-235: Running Cloudera Private Cloud

Empowering enterprise data

Course Overview

Course Type

Instructor-led training course

Level

Advanced

Duration

4 days

Platform

CDP Private Cloud

Topics Covered

- Architecture for CDP Data Services
- Docker Images and Containers
- Docker Private Registry
- Setup and install of Embeddec Containerized Services
- Configure Management Console and Environments
- Theory of Kubernetes
- Setup, Run, and Administer Cloudera Data Warehouse
- Install, Configure, and Administer Cloudera Data Engineering
- Install, Configure, and Administer Cloudera Machine Learning

About This Training

This four-day course teaches the architecture, deployment, configuration, and running of CDP Data Services on Embedded Containerized Services (ECS). CDP Data Services are state-of-the-art low code computing fusing together the entire data lifecycle into a single set of tools, reducing the costs of developing Use Cases while accelerating development and deployment.

The course begins with practices recommended for managing Docker images and containers resulting in the building of a Docker private registry. The Docker private registry is used to deploy the Data Services cluster on ECS. Students will learn to install, configure, and validate Cloudera Data Engineering, Cloudera Data Warehouse, and Cloudera Machine Learning. Exercises focus on learning Kubernetes, installing Private Cloud Embedded Container Service (ECS), and deploying Cloudera Data Services. The course includes requirements for networking and hardware, and explanations of Kubernetes pods dynamically scaling to support CDP Data Services.

Who Should Take This Course?

This immersion course is intended for CDP Administrators who are advancing into CDP Data Services running in a private cloud environment. We recommend a minimum of 3 to 5 years of system administration experience in industry. Students must have proficiency in Linux Command Line Interface, knowledge of Identity Management, Transport Layer Security, and Kerberos. Experience with SQL select statements is helpful. Prior experience with Cloudera products is expected, experience with CDP, CDH, or HDP is sufficient. Students must have access to the Internet to reach Amazon Web Services.

Recommended Prerequisite Courses

- *ADMIN-230 Operating Cloudera Data Platform* (Cloudera Educational Services OnDemand)
- *ADMIN-232 Building Secure Cloudera Clusters* (Cloudera Educational Services Instructor-led Training)

CDP Admin: Running Private Cloud

Training Outline (Page 2 of 3)

Big Data to Big Compute

- The importance of microservices
- Overview of containerized applications
- Design Principles for CDP Data Services

Docker Images and Containers

- Docker cli for images and containers
- Building a containerized application

Architecture for Cloudera Data Platform

- Architecture for Data Platform, Data Lake, and Data Services
- Architecture for Embedded Container Service

Docker Private Registry

- Build a Docker private registry
- Load Cloudera image packages

Clusters for Embedded Containerized Services

- Network requirements
- Hardware requirements
- Cluster recommendations

Shared Data Experience

- The Importance of Ranger, Atlas, and Hive Metastore
- Create Ranger policies on Hive databases in support of Cloudera Data Warehouse

Set Up Embedded Container Services

- Managing an air gapped repo
- Securing a Metastore database
- Creating wildcard DNS and wildcard certificates

Install Embedded Containerized Services

- Install Embedded Containerized Services
- Manage ECS Docker registry

Setup Management Console

- Configure secure LDAP
- Configure an administrator

Setup Environments

- Add users and groups
- Create environments
- Assign users to environments

Operate Management Console

- Manage alerts
- Create support bundles

Theory for Kubernetes

- Theory of Kubernetes clusters
- Theory of Kubernetes namespaces
- Theory of Kubernetes deployments
- Theory of Embedded Containerized Services

Manage Kubernetes

- Use Kubernetes web UI
- Use kubectl CLI
- Use K9 terminal interface

Manage Storage for ECS

- Architecture for Storage
- Architecture for Longhorn

Run Cloudera Data Warehouse

- Create virtual warehouses for Hive and Impala
- Administrator virtual warehouses

Run Cloudera Data Engineering

- Setup principals and wildcard certificates
- Create virtual clusters for Spark
- Administrator virtual clusters

Run Cloudera Data Engineering

- Create virtual clusters for Spark
- Administrator virtual clusters
- Administrator workspaces

Administration for Embedded Containerized Services

- Stop and start CDP clusters

CDP Admin: Running Private Cloud

Training Outline (Page 3 of 3)

- Uninstall ECS