



Course Data Sheet

Operations Management with the Hortonworks Data Platform

Overview

This 4-day course covers administration tasks for Hadoop 2.0 clusters/

Description

The course presents content related to the deployment lifecycle for a multi-node Hadoop cluster including: installation, configuration, monitoring, scaling and how Hadoop works with Big Data.

Duration

4 days

Price

For information about pricing, email us at sales-training@hortonworks.com

Prerequisites

Students should have a basic understanding of Hadoop and Linux environments.

Target Audience

This course is designed for IT administrators and operators responsible for installing, configuring and supporting an Apache Hadoop 2.0 deployment in a Linux environment.

Format

50% Instructor-led lecture/discussion
50% hands-on labs.

Course Objectives

After completing this course, students should be able to:

- Describe various tools and frameworks in the Hadoop 2.0 ecosystem
- Describe the Hadoop Distributed File System (HDFS) architecture
- Install and configure an HDP 2.0 cluster
- Use Ambari to monitor and manage a cluster
- Describe how files are written to and stored in HDFS
- Perform a file system check using command line and browser-based tools
- Configure the replication factor of a file
- Mount HDFS to a local filesystem using the NFS Gateway
- Deploy and configure YARN on a cluster
- Configure and troubleshoot MapReduce jobs
- Describe how YARN jobs are scheduled
- Configure the capacity and fair schedulers of the ResourceManager
- Use WebHDFS to access a cluster over HTTP
- Configure a Hiveserver
- Describe how Hive tables are created and populated
- Use Sqoop to transfer data between Hadoop and a relational database
- Use Flume to ingest streaming data into HDFS
- Deploy and run an Oozie workflow
- Commission and decommission worker nodes
- Configure a cluster to be rack-aware
- Implement and configure NameNode HA
- Secure a Hadoop cluster

Agenda

Day 1

- Unit 1: Introduction to HDP and Hadoop 2.0
- Unit 2: HDFS Architecture
- Unit 3: Installation Prerequisites and Planning • Unit 4: Configuring Hadoop
- Unit 5: Ensuring Data Integrity

Day 2

- Unit 6: HDFS NFS Gateway
- Unit 7: YARN Architecture and MapReduce
- Unit 8: Job Schedulers
- Unit 9: Enterprise Data Movement
- Unit 10: HDFS Web Services

Day 3

- Unit 11: Hive Administration
- Unit 12: Transferring Data with Sqoop
- Unit 13: Flume
- Unit 14: Oozie
- Unit 15: Monitoring HDP2 Services

Day 4

- Unit 16: Commissioning and Decommissioning Nodes
- Unit 17: Backup and Recovery
- Unit 18: Rack Awareness and Topology
- Unit 19: NameNode HA
- Unit 20: Securing HDP

Lab Content

Students will work through the following lab exercises using the Hortonworks Data Platform 2.0:

- Install HDP 2.0 using Ambari
- Add a new node to the cluster
- Stop and start HDP services
- Use HDFS commands
- Verify data with block scanner and fsck
- Mount HDFS to a local file system
- Troubleshoot a MapReduce job
- Configure the capacity scheduler
- Use `distcp` to copy data from a remote cluster
- Use WebHDFS
- Understanding Hive tables
- Use Sqoop to transfer data
- Install and test Flume
- Run an Oozie workflow
- Commission and decommission a worker node
- Use HDFS snapshots
- Configure rack awareness
- Implement NameNode HA
- Secure an HDP cluster

Additional Information

- For availability of individual seats in our open enrollment classes please visit us at www.hortonworks.com/training
- Please contact sales-training@hortonworks.com to discuss your specific training needs