



---

## Cloudera Administrator Training for Apache Hadoop

**Duration: 4 Days**    **Course Code: GK3901**

---

### Overview:

In this hands-on course, you will be introduced to the basics of Hadoop, Hadoop Distributed File System (HDFS), MapReduce, Hive, Pig, and HBase. You will cover core administration skills, such as cluster deployment, job management, and ongoing Hadoop maintenance and monitoring, as you gain the expertise to support your environments in day-to-day activities. This course covers concepts addressed on the Cloudera Certified Administrator for Apache Hadoop (CCAHA) exam and includes a CCAHA exam voucher you'll receive at the end of class.

---

### Target Audience:

System administrators looking to understand all of the steps necessary to operate and manage Apache Hadoop clusters

---

### Objectives:

- HDFS and MapReduce
  - Optimal hardware configurations for Hadoop clusters
  - Network considerations to take into account when building out your cluster
  - Configure Hadoop options for best cluster performance
  - Configure the FairScheduler to provide service-level agreements for multiple users of a cluster
  - Maintain and monitor your cluster
  - Load data into the cluster from dynamically generated files using Flume and from relational database management systems using Sqoop
  - System administration issues with other Hadoop projects such as Hive, Pig, and HBase
- 

### Prerequisites:

- Basic level of Linux system administration experience
- Prior knowledge of Apache Hadoop is not required

### Testing and Certification

**This course is part of the following programs or tracks:**

- CCAHA: Cloudera Certified Administrator for Apache Hadoop (CDH3)
- 

### Follow-on-Courses:

- Cloudera Training for Apache HBase
  - Cloudera Training for Apache Hive and Pig
-

## Content:

### Hadoop and HDFS

- Why Hadoop?
  - HDFS
  - MapReduce
  - Hive, Pig, HBase, and Other Ecosystem Projects
  - Choosing the Right Hardware
  - Node Topologies
  - Choosing the Right Software
  - Using SCM Express for Easy Installation
  - Typical Configuration Parameters
  - Configuring Rack Awareness
  - Using Configuration Management Tools
  - FIFO Scheduler
  - Fair Scheduler
  - Copying Data with Distcp
  - Rebalancing Cluster Nodes
  - Adding and Removing Cluster Nodes
  - Backup and Restore
  - Upgrading and Migrating
  - NameNode Metadata
  - Using the NameNode and JobTracker Web UIs
  - Interpreting Job Logs
  - Monitoring with Ganglia
  - Other Monitoring Tools
  - General Optimization Tips
  - Benchmarking Your Cluster
  - Using Flume
  - Best Practices for Data Ingestion
  - Pig
  - HBase
  - Metastore
- 
- HDFS
  - MapReduce
  - Hive, Pig, HBase, and Other Ecosystem Projects
  - Choosing the Right Hardware
  - Node Topologies
  - Choosing the Right Software
  - Using SCM Express for Easy Installation
  - Typical Configuration Parameters
  - Configuring Rack Awareness
  - Using Configuration Management Tools
  - FIFO Scheduler
  - Fair Scheduler
  - Copying Data with Distcp
  - Rebalancing Cluster Nodes
  - Adding and Removing Cluster Nodes
  - Backup and Restore
  - Upgrading and Migrating
  - NameNode Metadata
  - Using the NameNode and JobTracker Web UIs
  - Interpreting Job Logs
  - Monitoring with Ganglia
  - Other Monitoring Tools

### Managing and Scheduling Jobs

- Starting and Stopping MapReduce Jobs
  - HDFS
  - MapReduce
  - Hive, Pig, HBase, and Other Ecosystem Projects
  - Choosing the Right Hardware
  - Node Topologies
  - Choosing the Right Software
  - Using SCM Express for Easy Installation
  - Typical Configuration Parameters
  - Configuring Rack Awareness
  - Using Configuration Management Tools
  - FIFO Scheduler
  - Fair Scheduler
  - Copying Data with Distcp
  - Rebalancing Cluster Nodes
  - Adding and Removing Cluster Nodes
  - Backup and Restore
  - Upgrading and Migrating
  - NameNode Metadata
  - Using the NameNode and JobTracker Web UIs
  - Interpreting Job Logs
  - Monitoring with Ganglia
  - Other Monitoring Tools
  - General Optimization Tips
  - Benchmarking Your Cluster
  - Using Flume
  - Best Practices for Data Ingestion
  - Pig
  - HBase
  - Metastore
- 
- HDFS
  - MapReduce
  - Hive, Pig, HBase, and Other Ecosystem Projects
  - Choosing the Right Hardware
  - Node Topologies
  - Choosing the Right Software
  - Using SCM Express for Easy Installation
  - Typical Configuration Parameters
  - Configuring Rack Awareness
  - Using Configuration Management Tools
  - FIFO Scheduler
  - Fair Scheduler
  - Copying Data with Distcp
  - Rebalancing Cluster Nodes
  - Adding and Removing Cluster Nodes
  - Backup and Restore
  - Upgrading and Migrating
  - NameNode Metadata
  - Using the NameNode and JobTracker Web UIs
  - Interpreting Job Logs
  - Monitoring with Ganglia
  - Other Monitoring Tools

- HDFS
  - MapReduce
  - Hive, Pig, HBase, and Other Ecosystem Projects
  - Choosing the Right Hardware
  - Node Topologies
  - Choosing the Right Software
  - Using SCM Express for Easy Installation
  - Typical Configuration Parameters
  - Configuring Rack Awareness
  - Using Configuration Management Tools
  - FIFO Scheduler
  - Fair Scheduler
  - Copying Data with Distcp
  - Rebalancing Cluster Nodes
  - Adding and Removing Cluster Nodes
  - Backup and Restore
  - Upgrading and Migrating
  - NameNode Metadata
  - Using the NameNode and JobTracker Web UIs
  - Interpreting Job Logs
  - Monitoring with Ganglia
  - Other Monitoring Tools
  - General Optimization Tips
  - Benchmarking Your Cluster
  - Using Flume
  - Best Practices for Data Ingestion
  - Pig
  - HBase
  - Metastore
- 
- HDFS
  - MapReduce
  - Hive, Pig, HBase, and Other Ecosystem Projects
  - Choosing the Right Hardware
  - Node Topologies
  - Choosing the Right Software
  - Using SCM Express for Easy Installation
  - Typical Configuration Parameters
  - Configuring Rack Awareness
  - Using Configuration Management Tools
  - FIFO Scheduler
  - Fair Scheduler
  - Copying Data with Distcp
  - Rebalancing Cluster Nodes
  - Adding and Removing Cluster Nodes
  - Backup and Restore
  - Upgrading and Migrating
  - NameNode Metadata
  - Using the NameNode and JobTracker Web UIs
  - Interpreting Job Logs
  - Monitoring with Ganglia
  - Other Monitoring Tools
  - General Optimization Tips
  - Benchmarking Your Cluster
  - Using Flume
  - Best Practices for Data Ingestion

- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

#### Planning Your Hadoop Cluster

- General Planning Considerations

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating

- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

#### Cluster Maintenance

- Checking HDFS with Fsck
- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating

- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster

- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes

- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes

- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

#### Populating HDFS from External Sources

- Using Sqoop

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker

- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

#### Deploying Your Cluster

- Installing Hadoop

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness

- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp

- Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

#### Installing and Managing Other Hadoop Projects

- Hive

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler

- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation

- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

#### Cluster Monitoring, Troubleshooting, and Optimizing

- Hadoop Log Files

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software

- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- HDFS
- MapReduce
- Hive, Pig, HBase, and Other Ecosystem Projects
- Choosing the Right Hardware
- Node Topologies
- Choosing the Right Software
- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

#### Labs

- Install a Pseudo-Distributed Cluster
- Install a Hadoop Cluster
- Manage Jobs
- Use the FairScheduler
- Break the Cluster
- Verify the Cluster's Self-Healing Features
- Back Up and Restoring
- Configure the Hive Shared

- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

- Using SCM Express for Easy Installation
- Typical Configuration Parameters
- Configuring Rack Awareness
- Using Configuration Management Tools
- FIFO Scheduler
- Fair Scheduler
- Copying Data with Distcp
- Rebalancing Cluster Nodes
- Adding and Removing Cluster Nodes
- Backup and Restore
- Upgrading and Migrating
- NameNode Metadata
- Using the NameNode and JobTracker Web UIs
- Interpreting Job Logs
- Monitoring with Ganglia
- Other Monitoring Tools
- General Optimization Tips
- Benchmarking Your Cluster
- Using Flume
- Best Practices for Data Ingestion
- Pig
- HBase
- Metastore

---

### Further Information:

For More information, or to book your course, please call us on Head Office 01189 123456 / Northern Office 0113 242 5931

[info@globalknowledge.co.uk](mailto:info@globalknowledge.co.uk)

[www.globalknowledge.co.uk](http://www.globalknowledge.co.uk)

Global Knowledge, Mulberry Business Park, Fishponds Road, Wokingham Berkshire RG41 2GY UK