

Hadoop for Systems Administrators (GL660) H6C60S

HPE course number	H6C60S
Course length	3 days
Delivery mode	ILT, VILT
View schedule, local pricing, and register	View now
View related courses	View now

Why HPE Education Services?

- IDC MarketScape leader 4 years running for IT education and training*
- Recognized by IDC for leading with global coverage, unmatched technical expertise, and targeted education consulting services*
- Key partnerships with industry leaders OpenStack®, VMware®, Linux®, Microsoft®, ITIL, PMI, CSA, and (ISC)2
- Complete continuum of training delivery options—self-paced eLearning, custom education consulting, traditional classroom, video on-demand instruction, live virtual instructor-led with hands-on lab, dedicated onsite training
- Simplified purchase option with HPE Training Credits

This course covers the essentials of deploying and managing an Apache™ Hadoop® cluster. The course is lab intensive with each participant creating their own Hadoop cluster using either the CDH (Cloudera's Distribution, including Apache Hadoop) or Hortonworks Data Platform stacks. Core Hadoop services are explored in depth with emphasis on troubleshooting and recovering from common cluster failures. The fundamentals of related services such as Ambari, Zookeeper, Pig, Hive, HBase, Sqoop, Flume, and Oozie are also covered. The course is approximately 60% lecture and 40% labs.

Audience

Systems Administrators who will be responsible for managing and administering Hadoop clusters.

Prerequisites

Qualified participants should be comfortable with the Linux commands and have some systems administration experience, but do not need previous Hadoop experience.

Supported Distributions

Red Hat Enterprise Linux 7

Detailed Course Outline

Module 1: Hadoop Overview	<ul style="list-style-type: none"> • Data Analysis • Big Data • Origins of Hadoop • Hadoop Marketplace • Hadoop Core • Hadoop Ecosystem: • Hadoop Ecosystem (cont): • Hadoop Ecosystem (cont): 	<ul style="list-style-type: none"> • Hadoop Ecosystem (cont): • Hadoop Ecosystem (cont): • Cluster Architecture • Hardware/Software Requirements • Running Commands on Multiple Systems <p>Lab Tasks</p> <ul style="list-style-type: none"> • Running Commands on Multiple Hosts • Preparing to Install Hadoop
Module 2: HDFS	<ul style="list-style-type: none"> • Design Goals • Design • Blocks • Block Replication • Namenode Daemon • Secondary Namenode Daemon • Datanode Daemon • Accessing HDFS 	<ul style="list-style-type: none"> • Permissions and Users • Adding and Removing Datanodes • Balancing <p>Lab Tasks</p> <ul style="list-style-type: none"> • Single Node HDFS • Multi-node HDFS • Files and HDFS • Managing and Maintaining HDFS
Module 3: YARN	<ul style="list-style-type: none"> • YARN Design Goals • YARN Architecture • Resource Manager • Node Manager • Containers 	<ul style="list-style-type: none"> • YARN: Other Important Features • Slider <p>Lab Tasks</p> <ul style="list-style-type: none"> • YARN
Module 4: MapReduce	<ul style="list-style-type: none"> • MapReduce • Terminology and Data Flow 	<p>Lab Tasks</p> <ul style="list-style-type: none"> • Mapreduce
Module 5: Installing Hadoop with Ambari	<p>Lab Tasks</p> <ul style="list-style-type: none"> • CDH Uninstall 	<ul style="list-style-type: none"> • Installing Hadoop with Ambari • Tez
Module 6: Data Ingestion	<ul style="list-style-type: none"> • Sqoop • Flume • Kafka 	<p>Lab Tasks</p> <ul style="list-style-type: none"> • Sqoop
Module 7: Data Lineage and Governance	<ul style="list-style-type: none"> • Falcon • Atlas 	<ul style="list-style-type: none"> • Oozie
Module 8: Data Processing Frameworks	<ul style="list-style-type: none"> • The Bane of MapReduce • Tez overview • Pig • Hive • Spark 	<ul style="list-style-type: none"> • Storm • Solr • Solr (cont) • Lab Tasks • Pig
Module 9: NoSQL Implementations	<ul style="list-style-type: none"> • HBase • Phoenix • Module 10: Cluster Management 	<ul style="list-style-type: none"> • Ambari Metrics System (AMS) • Zookeeper

Learn more at
hpe.com/ww/learnbigdata

Follow us:



© Copyright 2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

c04577721, July 2017, H6C60S C.01