APACHE-HADOOP-DEVELOPER Q&As

Hadoop 2.0 Certification exam for Pig and Hive Developer

Pass Hortonworks APACHE-HADOOP-DEVELOPER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/apache-hadoop-developer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Hortonworks Official Exam Center

Instant Download After Purchase
100% Money Back Guarantee
365 Days Free Update
800,000+ Satisfied Customers

QUESTION 1

You need to perform statistical analysis in your MapReduce job and would like to call methods in the Apache Commons Math library, which is distributed as a 1.3 megabyte Java archive (JAR) file. Which is the best way to make this library available to your MapReducer job at runtime?

A. Have your system administrator copy the JAR to all nodes in the cluster and set its location in the HADOOP_CLASSPATH environment variable before you submit your job.

B. Have your system administrator place the JAR file on a Web server accessible to all cluster nodes and then set the HTTP_JAR_URL environment variable to its location.

C. When submitting the job on the command line, specify the ?ibjars option followed by the JAR file path.

D. Package your code and the Apache Commands Math library into a zip file named JobJar.zip

Correct Answer: C

Explanation: The usage of the jar command is like this,

Usage: hadoop jar [mainClass] args...

If you want the commons-math3.jar to be available for all the tasks you can do any one of these

1. Copy the jar file in $HADOOP_HOME/lib dir

2.

or

Use the generic option -libjars.

QUESTION 2

Identify the MapReduce v2 (MRv2 / YARN) daemon responsible for launching application containers and monitoring application resource usage?

A. ResourceManager

B. NodeManager

C. ApplicationMaster

D. ApplicationMasterService

E. TaskTracker

F. JobTracker

Correct Answer: B

Reference: Apache Hadoop YARN ?Conceptsand; Applications

QUESTION 3

Which one of the following classes would a Pig command use to store data in a table defined in HCatalog?

A. org.apache.hcatalog.pig.HCatOutputFormat

B. org.apache.hcatalog.pig.HCatStorer

C. No special class is needed for a Pig script to store data in an HCatalog table

D. Pig scripts cannot use an HCatalog table

Correct Answer: B

QUESTION 4

You want to run Hadoop jobs on your development workstation for testing before you submit them to your production cluster. Which mode of operation in Hadoop allows you to most closely simulate a production cluster while using a single machine?

A. Run all the nodes in your production cluster as virtual machines on your development workstation.

B. Run the hadoop command with the –jt local and the –fs file:///options.

C. Run the DataNode, TaskTracker, NameNode and JobTracker daemons on a single machine.

D. Run simldooop, the Apache open-source software for simulating Hadoop clusters.

Correct Answer: C

QUESTION 5

To use a lava user-defined function (UDF) with Pig what must you do?

A. Define an alias to shorten the function name

B. Pass arguments to the constructor of UDFs implementation class

C. Register the JAR file containing the UDF

D. Put the JAR file into the userandapos;s home folder in HDFS

Correct Answer: C

APACHE-HADOOP- APACHE-HADOOP- APACHE-HADOOP-

DEVELOPER VCE Dumps DEVELOPER Study Guide DEVELOPER Exam Questions