

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.pass2lead.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

Which one of the following statements describes the relationship between the NodeManager and the ApplicationMaster?

- A. The ApplicationMaster starts the NodeManager in a Container
- B. The NodeManager requests resources from the ApplicationMaster
- C. The ApplicationMaster starts the NodeManager outside of a Container
- D. The NodeManager creates an instance of the ApplicationMaster

Correct Answer: D

QUESTION 2

What does the following command do? `register andapos;:/piggyban):/pig-files.jarandapos;;`

- A. Invokes the user-defined functions contained in the jar file
- B. Assigns a name to a user-defined function or streaming command
- C. Transforms Pig user-defined functions into a format that Hive can accept
- D. Specifies the location of the JAR file containing the user-defined functions

Correct Answer: D

QUESTION 3

You need to run the same job many times with minor variations. Rather than hardcoding all job configuration options in your drive code, you've decided to have your Driver subclass `org.apache.hadoop.conf.Configured` and implement the `org.apache.hadoop.util.Tool` interface.

Identify which invocation correctly passes `mapred.job.name` with a value of Example to Hadoop?

- A. `hadoop "mapred.job.name=Example" MyDriver input output`
- B. `hadoop MyDriver mapred.job.name=Example input output`
- C. `hadoop MyDrive ? mapred.job.name=Example input output`
- D. `hadoop setproperty mapred.job.name=Example MyDriver input output`
- E. `hadoop setproperty ("mapred.job.name=Example") MyDriver input output`

Correct Answer: C

Explanation: Configure the property using the `-D key=value` notation:

-D mapred.job.name=\\My Job\\

You can list a whole bunch of options by calling the streaming jar with just the -info argument

Reference: Python hadoop streaming : Setting a job name

QUESTION 4

Which TWO of the following statements are true regarding Hive? Choose 2 answers A. Useful for data analysts familiar with SQL who need to do ad-hoc queries

- B. Offers real-time queries and row level updates
- C. Allows you to define a structure for your unstructured Big Data
- D. Is a relational database

Correct Answer: AC

QUESTION 5

What is a SequenceFile?

- A. A SequenceFile contains a binary encoding of an arbitrary number of homogeneous writable objects.
- B. A SequenceFile contains a binary encoding of an arbitrary number of heterogeneous writable objects.
- C. A SequenceFile contains a binary encoding of an arbitrary number of WritableComparable objects, in sorted order.
- D. A SequenceFile contains a binary encoding of an arbitrary number key-value pairs. Each key must be the same type. Each value must be same type.

Correct Answer: D

Explanation: SequenceFile is a flat file consisting of binary key/value pairs.

There are 3 different SequenceFile formats:

Uncompressed key/value records.

Record compressed key/value records - only \\values\\ are compressed here. Block compressed key/value records - both keys and values are collected in \\blocks\\ separately and compressed. The size of the \\block\\ is configurable.

Reference: <http://wiki.apache.org/hadoop/SequenceFile>

[APACHE-HADOOP-DEVELOPER Practice Test](#)

[APACHE-HADOOP-DEVELOPER Study Guide](#)

[APACHE-HADOOP-DEVELOPER Exam](#)

[Questions](#)