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 are developing a MapReduce job for sales reporting. The mapper will process input keys representing the year (IntWritable) and input values representing product indentifies (Text).

Indentify what determines the data types used by the Mapper for a given job.

A. The key and value types specified in the JobConf.setMapInputKeyClass and JobConf.setMapInputValuesClass methods

B. The data types specified in HADOOP_MAP_DATATYPES environment variable

C. The mapper-specification.xml file submitted with the job determine the mapper\\'s input key and value types.

D. The InputFormat used by the job determines the mapper\\'s input key and value types.

Correct Answer: D

Explanation: The input types fed to the mapper are controlled by the InputFormat used. The default input

format, "TextInputFormat," will load data in as (LongWritable, Text) pairs. The long value is the byte offset

of the line in the file. The Text object holds the string contents of the line of the file.

Note: The data types emitted by the reducer are identified by setOutputKeyClass()

andsetOutputValueClass(). The data types emitted by the reducer are identified by setOutputKeyClass()

and setOutputValueClass().

By default, it is assumed that these are the output types of the mapper as well. If this is not the case, the

methods setMapOutputKeyClass() and setMapOutputValueClass() methods of the JobConf class will

override these.

Reference: Yahoo! Hadoop Tutorial, THE DRIVER METHOD

QUESTION 2

You are developing a combiner that takes as input Text keys, IntWritable values, and emits Text keys, IntWritable values. Which interface should your class implement?

A. Combiner

B. Mapper

C. Reducer

D. Reducer

E. Combiner

Correct Answer: D

QUESTION 3

Which best describes what the map method accepts and emits?

A. It accepts a single key-value pair as input and emits a single key and list of corresponding values as output.

B. It accepts a single key-value pairs as input and can emit only one key-value pair as output.

C. It accepts a list key-value pairs as input and can emit only one key-value pair as output.

D. It accepts a single key-value pairs as input and can emit any number of key-value pair as output, including zero.

Correct Answer: D

Explanation: public class Mapper extends Object Maps input key/value pairs to a set of intermediate key/value pairs.

Maps are the individual tasks which transform input records into a intermediate records. The transformed intermediate records need not be of the same type as the input records. A given input pair may map to zero or many output pairs. Reference: org.apache.hadoop.mapreduce

Class Mapper

QUESTION 4

Which best describes how TextInputFormat processes input files and line breaks?

A. Input file splits may cross line breaks. A line that crosses file splits is read by the RecordReader of the split that contains the beginning of the broken line.

B. Input file splits may cross line breaks. A line that crosses file splits is read by the RecordReaders of both splits containing the broken line.

C. The input file is split exactly at the line breaks, so each RecordReader will read a series of complete lines.

D. Input file splits may cross line breaks. A line that crosses file splits is ignored.

E. Input file splits may cross line breaks. A line that crosses file splits is read by the RecordReader of the split that contains the end of the broken line.

Correct Answer: A

Reference: How Map and Reduce operations are actually carried out

---

QUESTION 5

Which one of the following statements regarding the components of YARN is FALSE?

A. A Container executes a specific task as assigned by the ApplicationMaster

B. The ResourceManager is responsible for scheduling and allocating resources

C. A client application submits a YARW job to the ResourceManager

D. The ResourceManager monitors and restarts any failed Containers

Correct Answer: D

| Latest APACHE-HADOOP-DEVELOPER Dumps | APACHE-HADOOP-DEVELOPER Study Guide | APACHE-HADOOP-DEVELOPER Exam Questions |