VCE & PDF
Pass4itSure.com

CCD-410^{Q&As}

Cloudera Certified Developer for Apache Hadoop (CCDH)

Pass Cloudera CCD-410 Exam with 100% Guarantee

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

https://www.pass4itsure.com/ccd-410.html

100% Passing Guarantee
100% Money Back Assurance

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

❖ **Instant Download** After Purchase

❖ **100% Money Back** Guarantee

❖ **365 Days** Free Update

❖ **800,000+** Satisfied Customers

QUESTION 1

For each input key-value pair, mappers can emit:

A. As many intermediate key-value pairs as designed. There are no restrictions on the types of those key-value pairs (i.e., they can be heterogeneous).

B. As many intermediate key-value pairs as designed, but they cannot be of the same type as the input key-value pair.

C. One intermediate key-value pair, of a different type.

D. One intermediate key-value pair, but of the same type.

E. As many intermediate key-value pairs as designed, as long as all the keys have the same types and all the values have the same type.

Correct Answer: E

Mapper maps input key/value pairs to a set of intermediate key/value pairs.

Maps are the individual tasks that transform input records into intermediate records. The transformed intermediate records do not need to be of the same type as the input records. A given input pair may map to zero or many output pairs.

Reference: Hadoop Map-Reduce Tutorial

QUESTION 2

In a MapReduce job with 500 map tasks, how many map task attempts will there be?

A. It depends on the number of reduces in the job.

B. Between 500 and 1000.

C. At most 500.

D. At least 500.

E. Exactly 500.

Correct Answer: D

Explanation: From Cloudera Training Course: Task attempt is a particular instance of an attempt to execute a task There will be at least as many task attempts as there are tasks If a task attempt fails, another will be started by the JobTracker Speculative execution can also result in more task attempts than completed tasks

QUESTION 3

You wrote a map function that throws a runtime exception when it encounters a control character in input data. The input supplied to your mapper contains twelve such characters totals, spread across five file splits. The first four file splits each have two control characters and the last split has four control characters.

Indentify the number of failed task attempts you can expect when you run the job with mapred.max.map.attempts set to 4:

A. You will have forty-eight failed task attempts

B. You will have seventeen failed task attempts

C. You will have five failed task attempts

D. You will have twelve failed task attempts

E. You will have twenty failed task attempts

Correct Answer: E

There will be four failed task attempts for each of the five file splits.

Note:

When the jobtracker is notified of a task attempt that has failed (by the tasktracker's heartbeat call), it will reschedule execution of the task. The jobtracker will try to avoid rescheduling the task on a tasktracker where it has previously failed. Furthermore, if a task fails four times (or more), it will not be retried further. This value is configurable: the maximum number of attempts to run a task is controlled by the `mapred.map.max.attempts` property for map tasks and `mapred.reduce.max.attempts` for reduce tasks. By default, if any task fails four times (or whatever the maximum number of attempts is configured to), the whole job fails.

QUESTION 4

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 D 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

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 5

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

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

[CCD-410 VCE Dumps](#)   [CCD-410 Practice Test](#)   [CCD-410 Exam Questions](#)