

Examcollection

<http://www.ipass4sure.com/examcollection.htm>

CCD-410

Cloudera

Cloudera Certified Developer for Apache Hadoop (CCDH)

The CCD-410 practice exam is written and formatted by Certified Senior IT Professionals working in today's prospering companies and data centers all over the world! The CCD-410 Practice Test covers all the exam topics and objectives and will prepare you for success quickly and efficiently. The CCD-410 exam is very challenging, but with our CCD-410 questions and answers practice exam, you can feel confident in obtaining your success on the CCD-410 exam on your FIRST TRY!

Cloudera CCD-410 Exam Features

- Detailed questions and answers for CCD-410 exam
- Try a demo before buying any Cloudera exam
- CCD-410 questions and answers, updated regularly
- Verified CCD-410 answers by Experts and bear almost 100% accuracy
- CCD-410 tested and verified before publishing
- CCD-410 examcollection vce questions with exhibits
- CCD-410 same questions as real exam with multiple choice options

Acquiring Cloudera certifications are becoming a huge task in the field of I.T. More over these exams like CCD-410 exam are now continuously updating and accepting this challenge is itself a task. This CCD-410 test is an important part of Cloudera certifications. We have the resources to prepare you for this. The CCD-410 exam is essential and core part of Cloudera certifications and once you clear the exam you will be able to solve the real life problems yourself. Want to take advantage of the Real CCD-410 Test and save time and money while developing your skills to pass your Cloudera CCD-410 Exam? Let us help you climb that ladder of success and pass your CCD-410 now!

DEMO EXAM

For Full Version visit

<http://www.ipass4sure.com/allexams.asp>

QUESTION: 1

When is the earliest point at which the reduce method of a given Reducer can be called?

- A. As soon as at least one mapper has finished processing its input split.
- B. As soon as a mapper has emitted at least one record.
- C. Not until all mappers have finished processing all records.
- D. It depends on the InputFormat used for the job.

Answer: C

Explanation:

In a MapReduce job reducers do not start executing the reduce method until the all Map jobs have completed. Reducers start copying intermediate key-value pairs from the mappers as soon as they are available. The programmer defined reduce method is called only after all the mappers have finished.

Note: The reduce phase has 3 steps: shuffle, sort, reduce. Shuffle is where the data is collected by the reducer from each mapper. This can happen while mappers are generating data since it is only a data transfer. On the other hand, sort and reduce can only start once all the mappers are done.

Why is starting the reducers early a good thing? Because it spreads out the data transfer from the mappers to the reducers over time, which is a good thing if your network is the bottleneck.

Why is starting the reducers early a bad thing? Because they "hog up" reduce slots while only copying data. Another job that starts later that will actually use the reduce slots now can't use them.

You can customize when the reducers startup by changing the default value of `mapred.reduce.slowstart.completed.maps` in `mapred-site.xml`. A value of 1.00 will wait for all the mappers to finish before starting the reducers. A value of 0.0 will start the reducers right away. A value of 0.5 will start the reducers when half of the mappers are complete. You can also change `mapred.reduce.slowstart.completed.maps` on a job-by-job basis.

Typically, keep `mapred.reduce.slowstart.completed.maps` above 0.9 if the system ever has multiple jobs running at once. This way the job doesn't hog up reducers when they aren't doing anything but copying data. If you only ever have one job running at a time, doing 0.1 would probably be appropriate.

Reference:

24 Interview Questions & Answers for Hadoop MapReduce developers, When is the reducers are started in a MapReduce job?

QUESTION: 2

Which describes how a client reads a file from HDFS?

- A. The client queries the NameNode for the block location(s). The NameNode returns the block location(s) to the client. The client reads the data directory off the DataNode(s).
- B. The client queries all DataNodes in parallel. The DataNode that contains the requested data responds directly to the client. The client reads the data directly off the DataNode.
- C. The client contacts the NameNode for the block location(s). The NameNode then queries the DataNodes for block locations. The DataNodes respond to the NameNode, and the NameNode redirects the client to the DataNode that holds the requested data block(s). The client then reads the data directly off the DataNode.
- D. The client contacts the NameNode for the block location(s). The NameNode contacts the DataNode that holds the requested data block. Data is transferred from the DataNode to the NameNode, and then from the NameNode to the client.

Answer: C

Explanation:

The Client communication to HDFS happens using Hadoop HDFS API. Client applications talk to the NameNode whenever they wish to locate a file, or when they want to add/copy/move/delete a file on HDFS. The NameNode responds the successful requests by returning a list of relevant DataNode servers where the data lives. Client applications can talk directly to a DataNode, once the NameNode has provided the location of the data.

Reference:

24 Interview Questions & Answers for Hadoop MapReduce developers,How the Client communicates with HDFS?

QUESTION: 3

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 <Text, IntWritable, Text, IntWritable>
- B. Mapper <Text, IntWritable, Text, IntWritable>
- C. Reducer <Text, Text, IntWritable, IntWritable>
- D. Reducer <Text, IntWritable, Text, IntWritable>
- E. Combiner <Text, Text, IntWritable, IntWritable>

Answer: D

QUESTION: 4

Identify the utility that allows you to create and run MapReduce jobs with any executable or script as the mapper and/or the reducer?

- A. Oozie
- B. Sqoop
- C. Flume
- D. Hadoop Streaming
- E. mapred

Answer: D

Explanation:

Hadoop streaming is a utility that comes with the Hadoop distribution. The utility allows you to create and run Map/Reduce jobs with any executable or script as the mapper and/or the reducer.

Reference:

<http://hadoop.apache.org/common/docs/r0.20.1/streaming.html>(Hadoop Streaming, second sentence)

QUESTION: 5

How are keys and values presented and passed to the reducers during a standard sort and shuffle phase of MapReduce?

- A. Keys are presented to reducer in sorted order; values for a given key are not sorted.
- B. Keys are presented to reducer in sorted order; values for a given key are sorted in ascending order.
- C. Keys are presented to a reducer in random order; values for a given key are not sorted.
- D. Keys are presented to a reducer in random order; values for a given key are sorted in ascending order.

Answer: A

Explanation:

Reducer has 3 primary phases:

1.Shuffle

The Reducer copies the sorted output from each Mapper using HTTP across the network.

2.Sort

The framework merge sorts Reducer inputs by keys (since different Mappers may have output the same key).

The shuffle and sort phases occur simultaneously i.e. while outputs are being fetched they are merged.

SecondarySort

To achieve a secondary sort on the values returned by the value iterator, the application should extend the key with the secondary key and define a grouping comparator. The

keys will be sorted using the entire key, but will be grouped using the grouping comparator to decide which keys and values are sent in the same call to reduce.

3.Reduce

In this phase the `reduce(Object, Iterable, Context)` method is called for each `<key, (collection of values)>` in the sorted inputs.

The output of the reduce task is typically written to a `RecordWriter` via `TaskInputOutputContext.write(Object, Object)`.

The output of the Reducer is not re-sorted.

Reference:

`org.apache.hadoop.mapreduce,Class`

`Reducer<KEYIN,VALUEIN,KEYOUT,VALUEOUT>`

QUESTION: 6

Assuming default settings, which best describes the order of data provided to a reducer's reduce method:

- A. The keys given to a reducer aren't in a predictable order, but the values associated with those keys always are.
- B. Both the keys and values passed to a reducer always appear in sorted order.
- C. Neither keys nor values are in any predictable order.
- D. The keys given to a reducer are in sorted order but the values associated with each key are in no predictable order

Answer: D

Explanation:

Reducer has 3 primary phases:

1.Shuffle

The Reducer copies the sorted output from each Mapper using HTTP across the network.

2.Sort

The framework merge sorts Reducer inputs by keys (since different Mappers may have output the same key).

The shuffle and sort phases occur simultaneously i.e. while outputs are being fetched they are merged.

SecondarySort

To achieve a secondary sort on the values returned by the value iterator, the application should extend the key with the secondary key and define a grouping comparator. The keys will be sorted using the entire key, but will be grouped using the grouping comparator to decide which keys and values are sent in the same call to reduce.

3.Reduce

In this phase the `reduce(Object, Iterable, Context)` method is called for each `<key, (collection of values)>` in the sorted inputs.

The output of the reduce task is typically written to a `RecordWriter` via `TaskInputOutputContext.write(Object, Object)`.

The output of the Reducer is not re-sorted.

Reference:

```
org.apache.hadoop.mapreduce,Class
Reducer<KEYIN,VALUEIN,KEYOUT,VALUEOUT>
```

QUESTION: 7

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. Identify 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

Answer: E

Explanation:

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: 8

You want to populate an associative array in order to perform a map-side join. You've decided to put this information in a text file, place that file into the DistributedCache and read it in your Mapper before any records are processed. Identify which method in the Mapper you should use to implement code for reading the file and populating the associative array?

- A. combine



Pass4sure Certification Exam Features;

- Pass4sure offers over **4500** Certification exams for professionals.
- More than **98,800** Satisfied Customers Worldwide.
- Average **99.8%** Success Rate.
- Over **150** Global Certification Vendors Covered.
- Services of **Professional & Certified Experts** available via support.
- Free **90 days** updates to match real exam scenarios.
- **Instant Download Access!** No Setup required.
- Price as low as **\$19**, which is 80% more **cost effective** than others.
- **Verified answers** researched by industry experts.
- Study Material **updated** on regular basis.
- Questions / Answers are downloadable in **PDF** format.
- Mobile Device Supported (**Android, iPhone, iPod, iPad**)
- **No authorization** code required to open exam.
- **Portable** anywhere.
- **Guaranteed Success.**
- **Fast**, helpful support **24x7**.



View list of All certification exams offered;
<http://www.ipass4sure.com/all exams.asp>

View list of All Study Guides (SG);
<http://www.ipass4sure.com/study-guides.asp>

View list of All Audio Exams (AE);
<http://www.ipass4sure.com/audio-exams.asp>

Download Any Certification Exam DEMO.
<http://www.ipass4sure.com/samples.asp>

To purchase Full version of exam click below;
<http://www.ipass4sure.com/all exams.asp>

3COM	CompTIA	Filemaker	IBM	LPI	OMG	Sun
ADOBE	ComputerAssociates	Fortinet	IISFA	McAfee	Oracle	Sybase
APC	CWNP	Foundry	Intel	McData	PMI	Symantec
Apple	DELL	Fujitsu	ISACA	Microsoft	Polycom	TeraData
BEA	ECCouncil	GuidanceSoftware	ISC2	Mile2	RedHat	TIA
BICSI	EMC	HDI	ISEB	NetworkAppliance	Sair	Tibco
CheckPoint	Enterasys	Hitachi	ISM	Network-General	SASInstitute	TruSecure
Cisco	ExamExpress	HP	Juniper	Nokia	SCP	Veritas
Citrix	Exin	Huawei	Legato	Nortel	See-Beyond	Vmware
CIW	ExtremeNetworks	Hyperion	Lotus	Novell	Google	

and many others.. See complete list [Here](#)

