



CCA-500^{Q&As}

Cloudera Certified Administrator for Apache Hadoop (CCA-H)

Pass Cloudera CCA-500 Exam with 100% Guarantee

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

<https://www.passapply.com/cca-500.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

A user comes to you, complaining that when she attempts to submit a Hadoop job, it fails. There is a Directory in HDFS named /data/input. The Jar is named j.jar, and the driver class is named DriverClass.

She runs the command:

```
Hadoop jar j.jar DriverClass /data/input/data/output
```

The error message returned includes the line: PrivilegedActionException as:training (auth:SIMPLE)
cause:org.apache.hadoop.mapreduce.lib.input.InvalidInputException: Input path does not exist: file:/data/input

What is the cause of the error?

- A. The user is not authorized to run the job on the cluster
- B. The output directory already exists
- C. The name of the driver has been spelled incorrectly on the command line
- D. The directory name is misspelled in HDFS
- E. The Hadoop configuration files on the client do not point to the cluster

Correct Answer: A

QUESTION 2

Your cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. What is the result when you execute: `hadoop jar SampleJar MyClass` on a client machine?

- A. SampleJar.Jar is sent to the ApplicationMaster which allocates a container for SampleJar.Jar
- B. Sample.jar is placed in a temporary directory in HDFS
- C. SampleJar.jar is sent directly to the ResourceManager
- D. SampleJar.jar is serialized into an XML file which is submitted to the ApplicationMaster

Correct Answer: A

QUESTION 3

Table schemas in Hive are:

- A. Stored as metadata on the NameNode
- B. Stored along with the data in HDFS
- C. Stored in the Metadata



D. Stored in ZooKeeper

Correct Answer: B

QUESTION 4

You observed that the number of spilled records from Map tasks far exceeds the number of map output records. Your child heap size is 1GB and your `io.sort.mb` value is set to 1000MB. How would you tune your `io.sort.mb` value to achieve maximum memory to disk I/O ratio?

A. For a 1GB child heap size an `io.sort.mb` of 128 MB will always maximize memory to disk I/O

B. Increase the `io.sort.mb` to 1GB

C. Decrease the `io.sort.mb` value to 0

D. Tune the `io.sort.mb` value until you observe that the number of spilled records equals (or is as close to equals) the number of map output records.

Correct Answer: D

QUESTION 5

Your cluster implements HDFS High Availability (HA). Your two NameNodes are named `nn01` and `nn02`. What occurs when you execute the command: `hdfs haadmin failover nn01 nn02`?

A. `nn02` is fenced, and `nn01` becomes the active NameNode

B. `nn01` is fenced, and `nn02` becomes the active NameNode

C. `nn01` becomes the standby NameNode and `nn02` becomes the active NameNode

D. `nn02` becomes the standby NameNode and `nn01` becomes the active NameNode

Correct Answer: B

[CCA-500 PDF Dumps](#)

[CCA-500 Practice Test](#)

[CCA-500 Study Guide](#)