


TestkingPass



Try Before You Buy

Download a free sample of any of our exam questions and answers

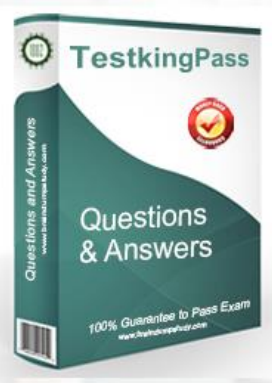
- ✓ 24/7 customer support, Secure shopping site
- ✓ Free One year updates to match real exam scenarios
- ✓ If you failed your exam after buying our products we will refund the full amount back to you.

Select a vendor...

Select an exam...

Your email address

[Free Download](#)



 HAPPY CUSTOMERS 51892	 DOWNLOADS 68912	 TEAM MEMBERS 56892	 SHARES 75162
-----------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------



<http://www.testkingpass.com>

Reliable test dumps & stable pass king & valid test questions

Exam : **CCA-505**

Title : Cloudera Certified
Administrator for Apache
Hadoop (CAHA) CDH5 Upgrade
Exam

Vendor : Cloudera

Version : DEMO

NO.1 During the execution of a MapReduce v2 (MRv2) job on YARN, where does the Mapper place the intermediate data each Map task?

- A. The Mapper stores the intermediate data on the node running the job's ApplicationMaster so that is available to YARN's ShuffleService before the data is presented to the Reducer
- B. The Mapper stores the intermediate data in HDFS on the node where the MAP tasks ran in the HDFS /usercache/&[user]sppcache/application_&(appid) directory for the user who ran the job
- C. YARN holds the intermediate data in the NodeManager's memory (a container) until it is transferred to the Reducers
- D. The Mapper stores the intermediate data on the underlying filesystem of the local disk in the directories yarn.nodemanager.local-dirs
- E. The Mapper transfers the intermediate data immediately to the Reducers as it generated by the Map task

Answer: D

NO.2 Which YARN process runs as "controller O" of a submitted job and is responsible for resource requests?

- A. ResourceManager
- B. NodeManager
- C. JobHistoryServer
- D. ApplicationMaster
- E. JobTracker
- F. ApplicationManager

Answer: D

NO.3 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 command:

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

The error message returned includes the line:

```
PrivilegedActionException as:training (auth:SIMPLE)
```

```
cause.apache.hadoop.mapreduce.lib.input.InvalidInputException: Input path does not exists: file  
:/data/input
```

What is the cause of the error?

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

Answer: A

NO.4 Given:

```
[user1@host1 ~] & yarn application -list
Total Applications: 3
```

Application-Id	Application-Name	Application-Type	User	Queue	State	Final State	Process	Tracking-URL
Application_1374638600275_0109	Sleep job	MAPREDUCE	user1	default	Killed	KILLED	100%	host1:54059
Application_1374638600275_0121	sleep job	MAPREDUCE	user1	default	SUCCEEDED	SUCCEEDED	100%	host1:19888/jobhistory/job/job_1374638600275_0121
Application_1374638600275_0020	sleep job	MAPREDUCE	user1	default	SUCCEEDED	SUCCEEDED	100%	host1:19888/jobhistory/job/job_1374638600275_0020

You want to clean up this list by removing jobs where the state is KILLED. What command you enter?

- A. Yarn application -kill application_1374638600275_0109
- B. Yarn radmin -refreshQueue
- C. Yarn application -refreshJobHistory
- D. Yarn radmin -kill application_1374638600275_0109

Answer: A

Reference: http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.1-latest/bk_usingapache-hadoop/content/common_mrv2_commands.html

NO.5 Your Hadoop cluster contains nodes in three racks. You have NOT configured the dfs.hosts property in the NameNode's configuration file. What results?

- A. No new nodes can be added to the cluster until you specify them in the dfs.hosts file
- B. Presented with a blank dfs.hosts property, the NameNode will permit DatNode specified in mapred.hosts to join the cluster
- C. Any machine running the DataNode daemon can immediately join the cluster
- D. The NameNode will update the dfs.hosts property to include machine running DataNode daemon on the next NameNode reboot or with the command dfsadmin -refreshNodes

Answer: C