CCA-505^{Q&As}

CCA-505^{Q&As}

Cloudera Certified Administrator for Apache Hadoop (CCAH) CDH5 Upgrade Exam

Pass Cloudera CCA-505 Exam with 100% Guarantee

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

https://www.passapply.com/cca-505.html

100% Passing Guarantee
100% Money Back Assurance

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

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 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 exits: 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

Correct Answer: A

QUESTION 2

Which Yarn daemon or service monitors a Container\\'s per-application resource usage (e.g, memory, CPU)?

A. NodeManager

B. ApplicationMaster

C. ApplicationManagerService

D. ResourceManager

Correct Answer: A

QUESTION 3

Your cluster is running MapReduce vserion 2 (MRv2) on YARN. Your ResourceManager is configured to use the FairScheduler. Now you want to configure your scheduler such that a new user on the cluster can submit jobs into their own queue application submission. Which configuration should you set?

A. You can specify new queue name when user submits a job and new queue can be created dynamically if yarn.scheduler.fair.user-as-default-queue = false

B. Yarn.scheduler.fair.user-as-default-queue = false and yarn.scheduler.fair.allow-undeclared- people = true

C. You can specify new queue name per application in allocation.fair.allow-undeclared-people = true automatically assigned to the application queue

D. You can specify new queue name when user submits a job and new queue can be created dynamically if the property yarn.scheduler.fair.allow-undecleared-pools = true

Correct Answer: A

QUESTION 4

You are migrating a cluster from MapReduce version 1 (MRv1) to MapReduce version2 (MRv2) on YARN. To want to maintain your MRv1 TaskTracker slot capacities when you migrate. What should you do?

A. Configure yarn.applicationmaster.resource.memory-mb and yarn.applicationmaster.cpu- vcores so that ApplicationMaster container allocations match the capacity you require.

B. You don\\'t need to configure or balance these properties in YARN as YARN dynamically balances resource management capabilities on your cluster

C. Configure yarn.nodemanager.resource.memory-mb and yarn.nodemanager.resource.cpu- vcores to match the capacity you require under YARN for each NodeManager

D. Configure mapred.tasktracker.map.tasks.maximum and mapred.tasktracker.reduce.tasks.maximum ub yarn.site.xml to match your cluster\\'s configured capacity set by yarn.scheduler.minimum-allocation

Correct Answer: C

QUESTION 5

You want to understand more about how users browse you public website. For example, you want to know which pages they visit prior to placing an order. You have a server farm of 200 web servers hosting your website. Which is the most efficient process to gather these web server logs into your Hadoop cluster for analysis?

A. Sample the web server logs web servers and copy them into HDFS using curl

B. Ingest the server web logs into HDFS using Flume

C. Import all users clicks from your OLTP databases into Hadoop using Sqoop

D. Write a MApReduce job with the web servers from mappers and the Hadoop cluster nodes reducers

E. Channel these clickstream into Hadoop using Hadoop Streaming

Correct Answer: AB

[Latest CCA-505 Dumps](#)   [CCA-505 PDF Dumps](#)   [CCA-505 VCE Dumps](#)