CCA-500^{Q&As}

Cloudera Certified Administrator for Apache Hadoop (CCAH)

Pass Cloudera CCA-500 Exam with 100% Guarantee

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

https://www.pass2lead.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

You want to understand more about how users browse your 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 across logs into your Hadoop cluster 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. Channel these clickstreams into Hadoop using Hadoop Streaming

D. Import all user clicks from your OLTP databases into Hadoop using Sqoop

E. Write a MapReeeduce job with the web servers for mappers and the Hadoop cluster nodes for reducers

Correct Answer: B

QUESTION 2

You are running a Hadoop cluster with a NameNode on host mynamenode. What are two ways to determine available HDFS space in your cluster?

A. Run hdfs fs du / and locate the DFS Remaining value

B. Run hdfs dfsadmin report and locate the DFS Remaining value

C. Run hdfs dfs / and subtract NDFS Used from configured Capacity

D. Connect to http://mynamenode:50070/dfshealth.jsp and locate the DFS remaining value

Correct Answer: B

QUESTION 3

Which YARN daemon or service negotiations map and reduce Containers from the Scheduler, tracking their status and monitoring progress?

A. NodeManager

B. ApplicationMaster

C. ApplicationManager

D. ResourceManager

Correct Answer: B

QUESTION 4

You use the hadoop fs put command to add a file "sales.txt" to HDFS. This file is small enough that it fits into a single block, which is replicated to three nodes in your cluster (with a replication factor of 3). One of the nodes holding this file (a single block) fails. How will the cluster handle the replication of file in this situation?

A. The file will remain under-replicated until the administrator brings that node back online

B. The cluster will re-replicate the file the next time the system administrator reboots the NameNode daemon (as long as the file\\'s replication factor doesn\\'t fall below)

C. This will be immediately re-replicated and all other HDFS operations on the cluster will halt until the cluster\\'s replication values are resorted

D. The file will be re-replicated automatically after the NameNode determines it is under- replicated based on the block reports it receives from the NameNodes

Correct Answer: D

QUESTION 5

A slave node in your cluster has 4 TB hard drives installed (4 x 2TB). The DataNode is configured to store HDFS blocks on all disks. You set the value of the dfs.datanode.du.reserved parameter to 100 GB. How does this alter HDFS block storage?

A. 25GB on each hard drive may not be used to store HDFS blocks

B. 100GB on each hard drive may not be used to store HDFS blocks

C. All hard drives may be used to store HDFS blocks as long as at least 100 GB in total is available on the node

D. A maximum if 100 GB on each hard drive may be used to store HDFS blocks

Correct Answer: C

[CCA-500 VCE Dumps](#)   [CCA-500 Practice Test](#)   [CCA-500 Exam Questions](#)