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.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

QUESTION 1

You suspect that your NameNode is incorrectly configured, and is swapping memory to disk. Which Linux commands help you to identify whether swapping is occurring? (Select all that apply)

A. free

B. df

C. memcat

D. top

E. jps

F. vmstat

G. swapinfo

Correct Answer: ADF

QUESTION 2

You are migrating a cluster from MApReduce version 1 (MRv1) to MapReduce version 2 (MRv2) on YARN. You 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.resource.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 mapred.tasktracker.map.tasks.maximum and mapred.tasktracker.reduce.tasks.maximum ub yarn-site.xml to match your cluster\\'s capacity set by the yarn-scheduler.minimum-allocation

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

Correct Answer: D

QUESTION 3

Assuming a cluster running HDFS, MapReduce version 2 (MRv2) on YARN with all settings at their default, what do you need to do when adding a new slave node to cluster?

A. Nothing, other than ensuring that the DNS (or/etc/hosts files on all machines) contains any entry for the new node.

B. Restart the NameNode and ResourceManager daemons and resubmit any running jobs.

C. Add a new entry to /etc/nodes on the NameNode host.

D. Restart the NameNode of dfs.number.of.nodes in hdfs-site.xml

Correct Answer: A

QUESTION 4

On a cluster running MapReduce v2 (MRv2) on YARN, a MapReduce job is given a directory of 10 plain text files as its input directory. Each file is made up of 3 HDFS blocks. How many Mappers will run?

A. We cannot say; the number of Mappers is determined by the ResourceManager

B. We cannot say; the number of Mappers is determined by the developer

C. 30

D. 3

E. 10

F. We cannot say; the number of mappers is determined by the ApplicationMaster

Correct Answer: E

QUESTION 5

Identify two features/issues that YARN is designated to address: (Choose two)

A. Standardize on a single MapReduce API

B. Single point of failure in the NameNode

C. Reduce complexity of the MapReduce APIs

D. Resource pressure on the JobTracker

E. Ability to run framework other than MapReduce, such as MPI

F. HDFS latency

Correct Answer: DE

QUESTION 6

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 ApplicatoionMaster

Correct Answer: A

QUESTION 7

On a cluster running CDH 5.0 or above, you use the hadoop fs put command to write a 300MB file into a previously empty directory using an HDFS block size of 64 MB. Just after this command has finished writing 200 MB of this file, what would another use see when they look in directory?

A. The directory will appear to be empty until the entire file write is completed on the cluster

B. They will see the file with a ._COPYING_ extension on its name. If they view the file, they will see contents of the file up to the last completed block (as each 64MB block is written, that block becomes available)

C. They will see the file with a ._COPYING_ extension on its name. If they attempt to view the file, they will get a ConcurrentFileAccessException until the entire file write is completed on the cluster

D. They will see the file with its original name. If they attempt to view the file, they will get a ConcurrentFileAccessException until the entire file write is completed on the cluster

Correct Answer: B

QUESTION 8

For each YARN job, the Hadoop framework generates task log file. Where are Hadoop task log files stored?

A. Cached by the NodeManager managing the job containers, then written to a log directory on the NameNode

B. Cached in the YARN container running the task, then copied into HDFS on job completion

C. In HDFS, in the directory of the user who generates the job

D. On the local disk of the slave mode running the task

Correct Answer: D

QUESTION 9

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

QUESTION 10

Cluster Summary:

45 files and directories, 12 blocks = 57 total. Heap size is 15.31 MB/193.38MB(7%)

Configured capacity	:	17.33GB
DFS Used	:	144KB
Non DFS Used	:	5.49GB
DFS Remaining	:	11.84GB
DFS Used %	:	0%
DFS Remaining %	:	68.32GB
Live Nodes	:	6
Dead Nodes	:	1
Decommissioning Nodes	:	0
Number of Under-Replicated Blocks	:	6

Refer to the above screenshot.

You configure a Hadoop cluster with seven DataNodes and on of your monitoring UIs displays the details

shown in the exhibit.

What does the this tell you?

A. The DataNode JVM on one host is not active

B. Because your under-replicated blocks count matches the Live Nodes, one node is dead, and your DFS Used % equals 0%, you can\\'t be certain that your cluster has all the data you\\'ve written it.

C. Your cluster has lost all HDFS data which had bocks stored on the dead DatNode

D. The HDFS cluster is in safe mode

Correct Answer: A