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.certbus.com/cca-505.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 a node to only swap Hadoop daemon data from RAM to disk when absolutely necessary. What should you do?

A. Delete the /swapfile file on the node

B. Set vm.swappiness to o in /etc/sysctl.conf

C. Set the ram.swap parameter to o in core-site.xml

D. Delete the /etc/swap file on the node

E. Delete the /dev/vmswap file on the node

Correct Answer: B

QUESTION 2

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 a cluster?

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

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

C. Increase the value of dfs.number.of.needs in hdfs-site.xml

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

E. Restart the NameNode daemon.

Correct Answer: B

QUESTION 3

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 4

What processes must you do if you are running a Hadoop cluster with a single NameNode and six DataNodes, and you want to change a configuration parameter so that it affects all six DataNodes.

A. You must modify the configuration file on each of the six DataNode machines.

B. You must restart the NameNode daemon to apply the changes to the cluster

C. You must restart all six DatNode daemon to apply the changes to the cluste

D. You don\\'t need to restart any daemon, as they will pick up changes automatically

E. You must modify the configuration files on the NameNode only. DataNodes read their configuration from the master nodes.

Correct Answer: BE

QUESTION 5

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 mode 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/and[user]sppcache/application_and(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

Correct Answer: D

QUESTION 6

For each YARN Job, the Hadoop framework generates task log files. Where are Hadoop\\'s files stored?

A. In HDFS, In the directory of the user who generates the job

B. On the local disk of the slave node running the task

C. Cached In the YARN container running the task, then copied into HDFS on fob completion

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

Correct Answer: B

QUESTION 7

Which is the default scheduler in YARN?

A. Fair Scheduler

B. FIFO Scheduler

C. Capacity Scheduler

D. YARN doesn\\'t configure a default scheduler. You must first assign a appropriate scheduler class in yarn-site.xml

Correct Answer: C

QUESTION 8

Your cluster has the following characteristics:

A rack aware topology is configured and on

Replication is not set to 3

Cluster block size is set to 64 MB

Which describes the file read process when a client application connects into the cluster and requests a 50MB file?

A. The client queries the NameNode which retrieves the block from the nearest DataNode to the client and then passes that block back to the client.

B. The client queries the NameNode for the locations of the block, and reads from a random location in the list it retrieves to eliminate network I/O leads by balancing which nodes it retrieves data from at any given time.

C. The client queries the NameNode for the locations of the block, and reads all three copies. The first copy to complete transfer to the client is the one the client reads as part of Hadoop\\'s

speculative execution framework.

D. The client queries the NameNode for the locations of the block, and reads from the first location in the list it receives.

Correct Answer: A

QUESTION 9

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 3)

A. free

B. df

C. memcat

D. top

E. vmstat

F. swapinfo

Correct Answer: ADE

QUESTION 10

You are upgrading a Hadoop cluster from HDFS and MapReduce version 1 (MRv1) to one running HDFS and MapReduce version 2 (MRv2) on YARN. You want to set and enforce a block of 128MB for all new files written to the cluster after the upgrade. What should you do?

A. Set dfs.block.size to 128M on all the worker nodes, on all client machines, and on the NameNode, and set the parameter to final.

B. Set dfs.block.size to 134217728 on all the worker nodes, on all client machines, and on the NameNode, and set the parameter to final.

C. Set dfs.block.size to 134217728 on all the worker nodes and client machines, and set the parameter to final. You do need to set this value on the NameNode.

D. Set dfs.block.size to 128M on all the worker nodes and client machines, and set the parameter to final. You do need to set this value on the NameNode.

E. You cannot enforce this, since client code can always override this value.

Correct Answer: C