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

Each node in your Hadoop cluster, running YARN, has 64 GB memory and 24 cores. Your yarn- site.xml

has the following configuration:

yarn.nodemanager.resource.memory-mb

32768

yarn.nodemanager.resource.cpu-vcores

23

You want YARN to launch no more than 16 containers per node. What should you do?

A. No action is needed: YARN\\'s dynamic resource allocation automatically optimizes the node memory and cores

B. Modify yarn-site.xml with the following property: yarn.nodemanager.resource.cpu-vcores 16

C. Modify yarn-site.xml with the following property: yarn.scheduler.minimum-allocation-mb 2048

D. Modify yarn-site.xml with the following property: yarn.scheduler.minimum-allocation-mb 4096

Correct Answer: B

QUESTION 2

You have a cluster running with the Fair Scheduler enabled. There are currently no jobs running on the cluster, and you submit a job A, so that only job A is running on the cluster. A while later, you submit Job B. now job A and Job B are running on the cluster at the same time. How will the Fair Scheduler handle these two jobs?

A. When job A gets submitted, it consumes all the tasks slots.

B. When job A gets submitted, it doesn\\'t consume all the task slots

C. When job B gets submitted, Job A has to finish first, before job B can scheduled

D. When job B gets submitted, it will get assigned tasks, while Job A continue to run with fewer tasks.

Correct Answer: C

QUESTION 3

You are configuring your cluster to run HDFS and MapReduce v2 (MRv2) on YARN. Which daemons need to be installed on your clusters master nodes? (Choose Two)

A. ResourceManager

B. DataNode

C. NameNode

D. JobTracker

E. TaskTracker

F. HMaster

Correct Answer: AC

QUESTION 4

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 5

Your cluster\\'s mapped-site.xml includes the following parameters

mapreduce.map.memory.mb

4096

mapreduce.reduce.memory,mb

8192

And your cluster\\'s yarn-site.xml includes the following parameters

yarn.nodemanager/vmen-pmem-ratio

2.1

What is the maximum amount of virtual memory allocated for each map before YARN will kill its Container?

A. 4 GB

B. 17.2 GB

C. 24.6 GB

D. 8.2 GB

Correct Answer: D

QUESTION 6

On a cluster running MapReduce v2 (MRv2) on YARN, a MapReduce job is given a directory of 10 plain text 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 RsourceManager

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

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

D. 30

E. 3

F. 10

Correct Answer: E

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

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 9

You are planning a Hadoop cluster and considering implementing 10 Gigabit Ethernet as the network fabric. Which workloads benefit the most from a faster network fabric?

A. When your workload generates a large amount of output data, significantly larger than amount of intermediate data

B. When your workload generates a large amount of intermediate data, on the order of the input data itself

C. When workload consumers a large amount of input data, relative to the entire capacity of HDFS

D. When your workload consists of processor-intensive tasks

Correct Answer: B

QUESTION 10

Your Hadoop cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. Can you configure a worker node to run a NodeManager daemon but not a DataNode daemon and still have a function cluster?

A. Yes. The daemon will receive data from the NameNode to run Map tasks

B. Yes. The daemon will get data from another (non-local) DataNode to run Map tasks

C. Yes. The daemon will receive Reduce tasks only

Correct Answer: A