VCE & PDF
GeekCert.com

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.geekcert.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

Which three basic configuration parameters must you set to migrate your cluster from MapReduce1 (MRv1) to MapReduce v2 (MRv2)?

A. Configure the NodeManager hostname and enable services on YARN by setting the following property in yarn-site.xml: yarn.nodemanager.hostname your_nodeManager_hostname

B. Configure the number of map tasks per job on YARN by setting the following property in mapredsite.xml: mapreduce.job.maps 2

C. Configure MapReduce as a framework running on YARN by setting the following property in mapredsite.xml: mapreduce.framework.name yarn

D. Configure the ResourceManager hostname and enable node services on YARN by setting the following property in yarn-site.xml: yarn.resourcemanager.hostname your_responseManager_hostname

E. Configure a default scheduler to run on YARN by setting the following property in sapred- site.xml: mapreduce.jobtracker.taskScheduler org.apache.hadoop.mapred.JobQueueTaskScheduler

F. Configure the NodeManager to enable MapReduce services on YARN by adding following property in yarn-site.xml: yarn.nodemanager.aux-services mapreduce_shuffle

Correct Answer: ABD

QUESTION 2

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 3

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 4

Which two are Features of Hadoop\\'s rack topology?

A. Configuration of rack awareness is accomplished using a configuration file. You cannot use a rack topology script.

B. Even for small clusters on a single rack, configuring rack awareness will improve performance.

C. Rack location is considered in the HDFS block placement policy

D. HDFS is rack aware but MapReduce daemons are not

E. Hadoop gives preference to Intra rack data transfer in order to conserve bandwidth

Correct Answer: BC

QUESTION 5

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

CCA-505 VCE Dumps   CCA-505 Study Guide   CCA-505 Braindumps