DUMPSARENA

Cloudera Certified Administrator for Apache Hadoop (CCAH)

Cloudera CCA-500

Version Demo

Total Demo Questions: 10

Total Premium Questions: 60

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which two features does Kerberos security add to a Hadoop cluster? (Choose two)

A. User authentication on all remote procedure calls (RPCs)

B. Encryption for data during transfer between the Mappers and Reducers

C. Encryption for data on disk (“at rest”)

D. Authentication for user access to the cluster against a central server

E. Root access to the cluster for users hdfs and mapred but non-root access for clients

ANSWER: A D

QUESTION NO: 2

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

ANSWER: B

Explanation:

failover – initiate a failover between two NameNodes

This subcommand causes a failover from the first provided NameNode to the second. If the first

NameNode is in the Standby state, this command simply transitions the second to the Active state without error. If the first NameNode is in the Active state, an attempt will be made to gracefully transition it to the Standby state. If this fails, the fencing methods (as configured by dfs.ha.fencing.methods) will be attempted in order until one of the methods succeeds. Only after this process will the second NameNode be transitioned to the Active state. If no fencing method succeeds, the second NameNode will not be transitioned to the Active state, and an error will be returned.

QUESTION NO: 3

You are configuring your cluster to run HDFS and MapReducer v2 (MRv2) on YARN. Which two daemons needs to be installed on your cluster’s master nodes? (Choose two)

A. HMaster

B. ResourceManager

C. TaskManager

D. JobTracker

E. NameNode

F. DataNode

ANSWER: B E

QUESTION NO: 4

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

ANSWER: B

QUESTION NO: 5

Your cluster is running MapReduce version 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 the property yarn.scheduler.fair.allow-undecleared-pools = true

B. Yarn.scheduler.fair.user.fair-as-default-queue = false and yarn.scheduler.fair.allow-undecleared-pools = true

C. You can specify new queue name when user submits a job and new queue can be created dynamically if yarn .schedule.fair.user-as-default-queue = false

D. You can specify new queue name per application in allocations.xml file and have new jobs automatically assigned to the application queue

ANSWER: A

QUESTION NO: 6

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

ANSWER: E

QUESTION NO: 7

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? (Choose two)

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

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

C. When Job A gets submitted, it doesn't consumes all the task slots.

D. When Job A gets submitted, it consumes all the task slots.

ANSWER: A D

QUESTION NO: 8

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

ANSWER: D E

QUESTION NO: 9

Choose three reasons why should you run the HDFS balancer periodically? (Choose three)

A. To ensure that there is capacity in HDFS for additional data

B. To ensure that all blocks in the cluster are 128MB in size

C. To help HDFS deliver consistent performance under heavy loads

D. To ensure that there is consistent disk utilization across the DataNodes

E. To improve data locality MapReduce

ANSWER: C D E

QUESTION NO: 10

Your cluster has the following characteristics:

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 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.

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

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

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

ANSWER: B