



Cloudera

CCA-505

*Cloudera Certified Administrator for Apache Hadoop
(CCA) CDH5 Upgrade*

Answer: B

QUESTION: 41

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 `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-undeclared-pools = true`

Answer: A

QUESTION: 42

A user comes to you, complaining that when she attempts to submit a Hadoop job, it fails. There is a directory in HDFS named `/data/input`. The Jar is named `j.jar`, and the driver class is named `DriverClass`. She runs command: `hadoop jar j.jar DriverClass /data/input/data/output` The error message returned includes the line: `PrivilegedActionException as:training (auth:SIMPLE) cause.apache.hadoop.mapreduce.lib.input.InvalidInputException: Input path does not exists: file ./data/input` What is the cause of the error?

- A. The Hadoop configuration files on the client do not point to the cluster
- B. The directory name is misspelled in HDFS
- C. The name of the driver has been spelled incorrectly on the command line
- D. The output directory already exists
- E. The user is not authorized to run the job on the cluster

Answer: A

QUESTION: 43

In CDH4 and later, which file contains a serialized form of all the directory and files inodes in the filesystem, giving the NameNode a persistent checkpoint of the filesystem metadata?

- A. fstime
- B. VERSION
- C. Fsimage_N (Where N reflects all transactions up to transaction ID N)
- D. Edits_N-M (Where N-M specifies transactions between transactions ID N and transaction ID N)

Answer: C

Reference:

<http://mikepluta.com/tag/namenode/>

QUESTION: 44

Which process instantiates user code, and executes map and reduce tasks on a cluster running MapReduce V2 (MRv2) on YARN?

- A. NodeManager
- B. ApplicationMaster
- C. ResourceManager
- D. TaskTracker
- E. JobTracker
- F. DataNode
- G. NameNode

Answer: E

QUESTION: 45

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-cores` 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-cores` to match the capacity you require under YARN for each NodeManager

D. Configure `mapred.tasktracker.map.tasks.maximum` and `mapred.tasktracker.reduce.tasks.maximum` in `yarn.site.xml` to match your cluster's configured capacity set by `yarn.scheduler.minimum-allocation`

Answer: C

Download Full Version From <https://www.certkillers.net>



DON'T KNOW
OR NO PREFERENCE

Pass your exam at First Attempt....Guaranteed!