

CCA-505^{Q&As}

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

Pass Cloudera CCA-505 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.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

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`

Correct Answer: A

QUESTION 2

Identify two features/issues that YARN is designed to address:

- A. Standardize on a single MapReduce API
- B. Single point of failure in the NameNode
- C. Reduce complexity of the MapReduce APIs
- D. Resource pressures on the JobTracker
- E. Ability to run frameworks other than MapReduce, such as MPI
- F. HDFS latency

Correct Answer: DE

QUESTION 3

You have installed a cluster running HDFS and MapReduce version 2 (MRv2) on YARN. You have no `dfs.hosts` entry(ies) in your `hdfs-site.xml` configuration file. You configure a new worker node by setting `fs.default.name` in its configuration files to point to the NameNode on your cluster, and you start the DataNode daemon on that worker node.

What do you have to do on the cluster to allow the worker node to join, and start storing HDFS blocks?

- A. Nothing; the worker node will automatically join the cluster when the DataNode daemon is started.
- B. Without creating a `dfs.hosts` file or making any entries, run the command `hadoop dfsadmin refreshHadoop` on the NameNode
- C. Create a `dfs.hosts` file on the NameNode, add the worker node's name to it, then issue the command `hadoop`

dfsadmin refreshNodes on the NameNode

D. Restart the NameNode

Correct Answer: B

QUESTION 4

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

Correct Answer: E

QUESTION 5

Which YARN daemon or service negotiates map and reduce Containers from the Scheduler, tracking their status and monitoring for progress?

A. ResourceManager

B. ApplicationMaster

C. NodeManager

D. ApplicationManager

Correct Answer: B

[CCA-505 Study Guide](#)

[CCA-505 Exam Questions](#)

[CCA-505 Braindumps](#)