



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.pass4itsure.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**

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 2

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 3**

Your Hadoop cluster contains nodes in three racks. You have NOT configured the `dfs.hosts` property in the `NameNode`'s configuration file. What results?

- A. No new nodes can be added to the cluster until you specify them in the `dfs.hosts` file
- B. Presented with a blank `dfs.hosts` property, the `NameNode` will permit `DatNode` specified in `mapred.hosts` to join the cluster
- C. Any machine running the `DataNode` daemon can immediately join the cluster
- D. The `NameNode` will update the `dfs.hosts` property to include machine running `DataNode` daemon on the next `NameNode` reboot or with the command `dfsadmin -refreshNodes`

Correct Answer: C

QUESTION 4

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/vmem-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 5

Assuming a cluster running HDFS, MapReduce version 2 (MRv2) on YARN with all settings at their default, what do you need to do when adding a new slave node to a cluster?

- A. Nothing, other than ensuring that DNS (or `/etc/hosts` files on all machines) contains an entry for the new node.



- B. Restart the NameNode and ResourceManager daemons and resubmit any running jobs
- C. Increase the value of dfs.number.of.replicas in hdfs-site.xml
- D. Add a new entry to /etc/nodes on the NameNode host.
- E. Restart the NameNode daemon.

Correct Answer: B

QUESTION 6

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 7

You are upgrading a Hadoop cluster from HDFS and MapReduce version 1 (MRv1) to one running HDFS and MapReduce version 2 (MRv2) on YARN. You want to set and enforce a block of 128MB for all new files written to the cluster after the upgrade. What should you do?

- A. Set dfs.block.size to 128M on all the worker nodes, on all client machines, and on the NameNode, and set the parameter to final.
- B. Set dfs.block.size to 134217728 on all the worker nodes, on all client machines, and on the NameNode, and set the parameter to final.
- C. Set dfs.block.size to 134217728 on all the worker nodes and client machines, and set the parameter to final. You do need to set this value on the NameNode.
- D. Set dfs.block.size to 128M on all the worker nodes and client machines, and set the parameter to final. You do need to set this value on the NameNode.
- E. You cannot enforce this, since client code can always override this value.

Correct Answer: C

QUESTION 8

Your cluster is configured with HDFS and MapReduce version 2 (MRv2) on YARN. What is the result when you



execute: `hadoop jar samplejar.jar MyClass` on a client machine?

- A. SampleJar.jar is sent to the ApplicationMaster which allocation a container for Sample.jar
- B. SampleJar.Jar is serialized into an XML file which is submitted to the ApplicationMaster
- C. SampleJar.Jar is sent directly to the ResourceManager
- D. SampleJar.Jar is placed in a temporary directly in HDFS

Correct Answer: A

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

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 `mapred-site.xml`: `mapreduce.job.maps 2`
- C. Configure MapReduce as a framework running on YARN by setting the following property in `mapred-site.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 `mapred-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



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/cca-505.html>

2024 Latest pass4itsure CCA-505 PDF and VCE dumps Download

[CCA-505 VCE Dumps](#)

[CCA-505 Practice Test](#)

[CCA-505 Exam Questions](#)