

KillTest

높은 품질, 더 나은 서비스



연습 문제

<https://www.killtest.kr>

1년간 무료 업데이트 제공

Exam : CCA-470

Title : Cloudera Certified
Administrator for Apache
Hadoop CDH4 Upgrade
Exam (CCA4)

Version : DEMO

1.Which two updates occur when a client application opens a stream to begin a file write on a cluster running MapReduce v1 (MRv1)?

- A.Once the write stream closes on the DataNode, the DataNode immediately initiates a block report to the NameNode.
- B.The change is written to the NameNode disk.
- C.The metadata in the RAM on the NameNode is flushed to disk.
- D.The metadata in RAM on the NameNode is flushed disk.
- E.The metadata in RAM on the NameNode is updated.
- F.The change is written to the edits file.

Answer: D, F

2.For a MapReduce job, on a cluster running MapReduce v1 (MRv1), what's the relationship between tasks and task templates?

- A.There are always at least as many task attempts as there are tasks.
- B.There are always at most as many tasks attempts as there are tasks.
- C.There are always exactly as many task attempts as there are tasks.
- D.The developer sets the number of task attempts on job submission.

Answer: C

3.What action occurs automatically on a cluster when a DataNode is marked as dead?

- A.The NameNode forces re-replication of all the blocks which were stored on the dead DataNode.
- B.The next time a client submits job that requires blocks from the dead DataNode, the JobTracker receives no heart beats from the DataNode.The JobTracker tells the NameNode that the DataNode is dead, which triggers block re-replication on the cluster.
- C.The replication factor of the files which had blocks stored on the dead DataNode is temporarily reduced, until the dead DataNode is recovered and returned to the cluster.
- D.The NameNode informs the client which write the blocks that are no longer available; the client then re-writes the blocks to a different DataNode.

Answer: A

4.How does the NameNode know DataNodes are available on a cluster running MapReduce v1 (MRv1)

- A.DataNodes listed in the dfs.hosts file.The NameNode uses as the definitive list of available DataNodes.
- B.DataNodes heartbeat in the master on a regular basis.
- C.The NameNode broadcasts a heartbeat on the network on a regular basis, and DataNodes respond.
- D.The NameNode send a broadcast across the network when it first starts, and DataNodes respond.

Answer: B

5.Which three distcp features can you utilize on a Hadoop cluster?

- A.Use distcp to copy files only between two clusters or more.You cannot use distcp to copy data between directories inside the same cluster.
- B.Use distcp to copy HBase table files.
- C.Use distcp to copy physical blocks from the source to the target destination in your cluster.
- D.Use distcp to copy data between directories inside the same cluster.
- E.Use distcp to run an internal MapReduce job to copy files.

Answer: B, D, E

6.How does HDFS Federation help HDFS Scale horizontally?

- A.HDFS Federation improves the resiliency of HDFS in the face of network issues by removing the NameNode as a single-point-of-failure.
- B.HDFS Federation allows the Standby NameNode to automatically resume the services of an active NameNode.
- C.HDFS Federation provides cross-data center (non-local) support for HDFS, allowing a cluster administrator to split the Block Storage outside the local cluster.
- D.HDFS Federation reduces the load on any single NameNode by using the multiple, independent NameNode to manage individual parts of the filesystem namespace.

Answer: D

7.Choose which best describe a Hadoop cluster's block size storage parameters once you set the HDFS default block size to 64MB?

- A.The block size of files in the cluster can be determined as the block is written.
- B.The block size of files in the Cluster will all be multiples of 64MB.
- C.The block size of files in the cluster will all at least be 64MB.
- D.The block size of files in the cluster will all be the exactly 64MB.

Answer: D

8.Which MapReduce daemon instantiates user code, and executes map and reduce tasks on a cluster running MapReduce v1 (MRv1)?

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

Answer: D

9.What two processes must you do if you are running a Hadoop cluster with a single NameNode and six DataNodes, and you want to change a configuration parameter so that it affects all six DataNodes.

- A.You must restart the NameNode daemon to apply the changes to the cluster
- B.You must restart all six DataNode daemons to apply the changes to the cluster.
- C.You don't need to restart any daemon, as they will pick up changes automatically.
- D.You must modify the configuration files on each of the six DataNode machines.
- E.You must modify the configuration files on only one of the DataNode machine
- F.You must modify the configuration files on the NameNode only.DataNodes read their configuration from the master nodes.

Answer: A, F

10.Identify the function performed by the Secondary NameNode daemon on a cluster configured to run

with a single NameNode.

A.In this configuration, the Secondary NameNode performs a checkpoint operation on the files by the NameNode.

B.In this configuration, the Secondary NameNode is standby NameNode, ready to failover and provide high availability.

C.In this configuration, the Secondary NameNode performs deal-time backups of the NameNode.

D.In this configuration, the Secondary NameNode servers as alternate data channel for clients to reach HDFS, should the NameNode become too busy.

Answer: A