
Question: 1

Which Hadoop component is responsible for managing the distributed file system metadata?

- A. NameNode
- B. Metanode
- C. DataNode
- D. NameSpaceManager

Answer:A

Question: 2

A NameNode in Hadoop 2.2 manages _____.

- A. Two namespaces: an active namespace and a backup namespace
- B. A single namespace
- C. An arbitrary number of namespaces
- D. No namespaces

Answer:B

Question: 3

Which Two of the following statements are true about hdfs? Choose 2 answers

- A. An HDFS file that is larger than dfs.block.size is split into blocks
- B. Blocks are replicated to multiple datanodes
- C. HDFS works best when storing a large number of relatively small files
- D. Block sizes for all files must be the same size

Answer:A, B

Question: 4

In Hadoop 2.2, which one of the following statements is true about a standby NameNode?
The Standby NameNode:

- A. Communicates directly with the active NameNode to maintain the state of the active NameNode.
- B. Receives the same block reports as the active NameNode.
- C. Runs on the same machine and shares the memory of the active NameNode.
- D. Processes all client requests and block reports from the appropriate DataNodes.

Answer:B

Question: 5

Which HDFS command uploads a local file X into an existing HDFS directory Y?

- A. `hadoop scp X Y`
- B. `hadoop fs -localPut X Y`
- C. `hadoop fs-put X Y`
- D. `hadoop fs -get X Y`

Answer:C

Question: 6

What does the following WebHDFS command do?
`Curl -1 -L "http://host:port/webhdfs/v1/foo/bar?op=OPEN"`

- A. Make a directory `/foo/bar`
- B. Read a file `/foo/bar`
- C. List a directory `/foo`
- D. Delete a directory `/foo/bar`

Answer:B

Question: 7

In Hadoop 2.2, which TWO of the following processes work together to provide automatic failover of the NameNode? Choose 2 answers

- A. ZKFailoverController
- B. ZooKeeper
- C. QuorumManager
- D. JournalNode

Answer:A, D

Question: 8

Which one of the following statements is FALSE regarding the communication between DataNodes and a federation of NameNodes in Hadoop 2.2?

- A. Each DataNode receives commands from one designated master NameNode.
- B. DataNodes send periodic heartbeats to all the NameNodes.
- C. Each DataNode registers with all the NameNodes.
- D. DataNodes send periodic block reports to all the NameNodes.

Answer:A

Question: 9

What is the term for the process of moving map outputs to the reducers?

- A. Reducing
- B. Combining
- C. Partitioning
- D. Shuffling and sorting

Answer:D

Question: 10

Which one of the following statements is true regarding a MapReduce job?

- A. The job's Partitioner shuffles and sorts all (key.value) pairs and sends the output to all reducers
- B. The default Hash Partitioner sends key value pairs with the same key to the same Reducer
- C. The reduce method is invoked once for each unique value
- D. The Mapper must sort its output of (key.value) pairs in descending order based on value

Answer:A
