


HADOOP MOCK TEST

<http://www.tutorialspoint.com>

Copyright © tutorialspoint.com

This section presents you various set of Mock Tests related to **Hadoop Framework**. You can download these sample mock tests at your local machine and solve offline at your convenience. Every mock test is supplied with a mock test key to let you verify the final score and grade yourself.



HADOOP MOCK TEST

Q 1 - The purpose of checkpoint node in a Hadoop cluster is to

- A - Check if the namenode is active
- B - Check if the fsimage file is in sync between namenode and secondary namenode
- C - Merges the fsimage and edit log and uploads it back to active namenode.
- D - Check which data nodes are unreachable.

Q 2 - When a backup node is used in a cluster there is no need of

- A - Check point node
- B - Secondary name node
- C - Secondary data node
- D - Rack awareness

Q 3 - Rack awareness in name node means

- A - It is aware how many racks are available in the cluster
- B - It is aware of the mapping between the node and the rack
- C - It is aware of the number of nodes in each of the rack
- D - It is aware which data nodes are unavailable in the cluster.

Q 4 - When a machine is declared as a datanode, the disk space in it

- A - Can be used only for HDFS storage
- B - Can be used for both HDFS and non-HDFS storage

C - Cannot be accessed by non-hadoop commands

D - cannot store text files.

Q 5 - When a file in HDFS is deleted by a user

A - it is lost forever

B - It goes to trash if configured.

C - It becomes hidden from the user but stays in the file system

D - File sin HDFS cannot be deleted

Q 6 - The source of HDFS architecture in Hadoop originated as

A - Google distributed filesystem

B - Yahoo distributed filesystem

C - Facebook distributed filesystem

D - Azure distributed filesystem

Q 7 - The inter process communication between different nodes in Hadoop uses

A - REST API

B - RPC

C - RMI

D - IP Exchange

Q 8 - The type of data Hadoop can deal with is

A - Structred

B - Semi-structured

C - Unstructured

D - All of the above

Q 9 - YARN stands for

A - Yahoo's another resource name

B - Yet another resource negotiator

C - Yahoo's archived Resource names

D - Yet another resource need.

Q 10 - The fully distributed mode of installation *without virtualization* needs a minimum of

- A - 2 physical mashines
- B - 3 Physical machines
- C - 4 Physical machines
- D - 1 Physical machine

Q 11 - Running Start-dfs.sh results in

- A - Starting namenode and datanode
- B - Starting namenode only
- C - Starting datanode only
- D - Starting namenode and resource manager

Q 12 - Which of the following is not a goal of HDFS?

- A - Fault detection and recovery
- B - Handle huge dataset
- C - Prevent deletion of data
- D - Provide high network bandwidth for data movement

Q 13 - The command "hadoop fs -test -z URI " gives the result 0 if

- A - if the path is a directory
- B - if the path is a file
- C - if the path is not empty
- D - if the file is zero length

Q 14 - In HDFS the files cannot be

- A - read
- B - deleted
- C - executed
- D - Archived

Q 15 - hadoop fs -expunge

- A - Gives the list of datanodes
- B - Used to delete a file
- C - Used to exchange a file between two datanodes.
- D - Empties the trash.

Q 16 - All the files in a directory in HDFS can be merged together using

- A - getmerge
- B - putmerge
- C - remerge
- D - mergeall

Q 17 - The replication factor of a file in HDFS can be changed using

- A - changerep
- B - rerep
- C - setrep
- D - xrep

Q 18 - The command used to copy a directory from one node to another in HDFS is

- A - rcp
- B - dcp
- C - drcp
- D - distcp

Q 19 - The archive file created in Hadoop always has the extension of

- A - .hrc
- B - .har
- C - .hrh
- D - .hrr

Q 20 - To unarchive an already archived file in Hadoop use the command

- A - unrar
- B - unhar
- C - cp
- D - cphar

Q 21 - The data from a remote Hadoop cluster can

- A - not be read by another Hadoop cluster
- B - be read using http
- C - be read using hhttp

D - be read using hftp

Q 22 - The purpose of starting namenode in the recovery mode is to

- A - Recover a failed namenode
- B - Recover a failed datanode
- C - Recover data from one of the metadata storage locations
- D - Recover data when there is only one metadata storage location

Q 23 - When you increase the number of files stored in HDFS, The memory required by namenode

- A - Increases
- B - Decreases
- C - Remains unchanged
- D - May increase or decrease

Q 24 - If we increase the size of files stored in HDFS without increasing the number of files, then the memory required by namenode

- A - Decreases
- B - Increases
- C - Remains unchanged
- D - May or may not increase

Q 25 - The current limiting factor to the size of a hadoop cluster is

- A - Excess heat generated in data center
- B - Upper limit of the network bandwidth
- C - Upper limit of the RAM in namenode
- D - 4000 data nodes

Q 26 - The decommission feature in hadoop is used for

- A - Decommissioning the namenode
- B - Decommissioning the data nodes
- C - Decommissioning the secondary namenode.
- D - Decommissioning the entire Hadoop cluster.

Q 27 - You can reserve the amount of disk usage in a data node by configuring the `dfs.datanode.du.reserved` in which of the following file

- A - Hdfs-site.xml
- B - Hdfs-defaukt.xml
- C - Core-site.xml
- D - Mapred-site.xml

Q 28 - The namenode loses its only copy of fsimage file. We can recover this from

- A - Datanodes
- B - Secondary namenode
- C - Checkpoint node
- D - Never

Q 29 - In a HDFS system with block size 64MB we store a file which is less than 64MB. Which of the following is true?

- A - The file will consume 64MB
- B - The file will consume more than 64MB
- C - The file will consume less than 64MB.
- D - Can not be predicted.

Q 30 - A running job in hadoop can

- A - Be killed with a command
- B - Can never be killed with a command
- C - Can be killed only by shutting down the name node
- D - Be paused and run again

Q 31 - The number of tasks a task tracker can accept depends on

- A - Maximum memory available in the node
- B - Not limited
- C - Number of slots configured in it
- D - As decided by the jobTracker

Q 32 - When a jobTracker schedules a task is first looks for

- A - A node with empty slot in the same rack as datanode
- B - Any node on the same rack as the datanode
- C - Any node on the rack adjacent to rack of the datanode
- D - Just any node in the cluster

Q 33 - The heartbeat signal are sent from

- A - JObtracker to Tasktracker
- B - Tasktracker to Job tracker
- C - Jobtracker to namenode
- D - Tasktracker to namenode

ANSWER SHEET

Question Number	Answer Key
-----------------	------------

1	C
2	A
3	B
4	B
5	B
6	A
7	B
8	D
9	B
10	A
11	A
12	C
13	D
14	C
15	D
16	A
17	C
18	D
19	B
20	C
21	D
22	D
23	A
24	A
25	C

26	B
27	A
28	C
29	C
30	A
31	C
32	A
33	B

Loading [MathJax]/jax/output/HTML-CSS/jax.js