

HBASE MOCK TEST

<http://www.tutorialspoint.com>

Copyright © tutorialspoint.com

This section presents you various set of Mock Tests related to **HBase**. 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.



HBASE MOCK TEST II

Q 1 - The data in a cell in Hbase table is identified using the four coordinates. Three of which are – rowkey, column family and column qualifier. The fourth coordinate used to identify each value in a cell is

- A - Sequence number
- B - Version number
- C - Serial number
- D - table name

Q 2 - Retrieving a batch of rows in every RPC call made by an API to a HBase database is called a

- A - Batch
- B - Scan
- C - Bulkrow
- D - Grouprow

Q 3 - A scan returns bulk of rows. But only a selected few rows can be fetched form a scan using a

- A - Group by clause
- B - Minimize clause
- C - Subset clause
- D - Filter clause

Q 4 - Filters in Hbase can be applied to

- A - Rowkeys
- B - Column qualifiers
- C - Data values
- D - All of the above

Q 5 - The command which allows you to change an integer value stored in Hbase cell without reading it first is

- A - Incrementcolumnvalue
- B - Incrementinteger
- C - Incrmentcellval
- D - Incrementnext

Q 6 - The number of columns Hbase table can hold is

- A - 256
- B - Depends on what value is configured
- C - No limit
- D - Depends on the HDFS block size

Q 7 - A small chunk of data residing in one machine which is part of a cluster of machines holding one Hbase table is known as

- A - Split
- B - Region
- C - Rowarea
- D - Tablearea

Q 8 - Servers that host regions of a Hbase table are called

- A - RegionServers
- B - Regional servers
- C - Hbase Servers
- D - Splitservers

Q 9 - Typically a Hbase Regionserver is collocated with

- A - HDFS Namenode
- B - HDFS datanode
- C - As a client to HDFS server

D - Tasktrackers

Q 10 - The size of a individual region is governed by the parameter

A - Hbase.region.size

B - Hbase.region.filesize

C - Hbase.region.max.filesize

D - Hbase.max.region.size

Q 11 - When a region becomes bigger in size, it

A - Gets trimmed

B - Spills into new machines

C - Is discarded

D - Is split into smaller regions

Q 12 - The two tables which are used to find where regions of various tables are hosted are

A - Regiontab and Metatab

B - Regionbase and Metabase

C - -ROOT- and .META.

D - -ROOT- and .REGION.

Q 13 - When a client application wants to access a row in a Hbase table it first queries the table

A - -ROOT-

B - .META.

C - .REGIONS.

D - .ALLREGIONS.

Q 14 - In any mapreduce Job Hbase can be used as a

A - Metadata store

B - Data source

C - Datanode

D - Metadata node

Q 15 - All MapReduce jobs reading from an Hbase table accept their [K1,V1] pair in the form of

A - [rowid:cell value]

B - [rowkey:scan result]

C - [column Family:cell value]

D - [column attribute:scan result]

Q 16 - When a map tasks in a mapreduce job reads from the Hbase table, it reads from

A - One row

B - One column family

C - One column

D - One region

Q 17 - The part of a Mapreduce Task which writes to a Hbase table is

A - Map

B - Reduce

C - Keys

D - none

Q 18 - While writing to Hbase using the Mapreduce tasks, each reduce tasks writes to

A - One region

B - Two regions

C - All the relevant regions

D - No regions

Q 19 - In a reduce-side join the Mapreduce step which is used to collocate the relevant records form the two joining data sets is

A - Map step

B - Reduce step

C - Shuffle and sort step

D - Final output step

Q 20 - The dis-advantage of reduce-side join is

A - Larger I/O costs because of shuffle and sort

B - Can process only small files for join

C - Cannot handle date data type while joining

D - May not give correct result always.

Q 21 - In a Map-Side join, we take rows from one table and map it with rows from the other table. The size of one of the table should be

- A - Enough to fit into memory
- B - Half the size of the other table
- C - Double the size of the other table
- D - Small enough to be located in one physical machine

Q 22 - Hbase stores data in

- A - A single filesystem available to all RegionServers
- B - As many filesystems as the number of regionServers
- C - One filesystem per column family
- D - One filesystem per table.

Q 23 - The number of namespaces, HDFS provides to the regionservers of a Hbase database is

- A - Equal to number of regionserver
- B - Half of the number of regionserver
- C - Double the number of regionserver
- D - One

Q 24 - Inserting data from one table to another in Hbase is

- A - Allowed
- B - Not allowed
- C - Possible using a third staging table
- D - Only one column family can be inserted at a time.

Q 25 - In Hbase we can create index on

- A - Multiple columns
- B - Key column
- C - Only a pair of columns
- D - Hbase does not support index creation.

ANSWER SHEET

Question Number Answer Key

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