

HIVE MOCK TEST

<http://www.tutorialspoint.com>

Copyright © tutorialspoint.com

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



HIVE MOCK TEST III

Q 1 - In case of one large table and 2 small tables, for an optimized query performance

- A - The largest one should be cached to memory and small ones should be streamed
- B - The small Ones should be cached and large one should be streamed
- C - All of the table should be cached
- D - All the tables should be streamed.

Q 2 - The DISTRIBUTED BY clause in hive

- A - comes Before the sort by clause
- B - comes after the sort by clause
- C - does not depend on position of sort by clause
- D - cannot be present along with the sort by clause

Q 3 - The DISTRIBUTED by clause is used to ensure that

- A - similar values go to the same reducer
- B - similar values go to same mapper
- C - similar values are distributed evenly across all the reducers
- D - similar values are distributed evenly across all the mapers

Q 4 - A view in Hive can be seen by using

- A - SHOW TABLES
- B - SHOW VIEWS

C - DESCRIBE VIEWS

D - VIEW VIEWS

Q 5 - A View in Hive can be dropped by using

A - DROP TABLE

B - DROP VIEW

C - DELETE VIEW

D - REMOVE VIEW

Q 6 - The name of a view in Hive

A - can be same as the name of another table in the same database

B - cannot be same as the name of another table in the same database

C - cannot contain a number

D - cannot be more than 10 character long

Q 7 - The query

```
Create table TABLE_NAME LIKE VIEW_NAME
```

A - creates a table which is copy of the view

B - is invalid

C - runs only if the view has data

D - runs only if the view is in same directory as the table

Q 8 - what can be altered about a view

A - its name

B - its location

C - its TBLPROPERTIES

D - The query it is based on

Q 9 - Which kind of keysCONSTRAINTS Hive can have?

A - Primary Keys

B - Foreign Keys

C - Unique Keys

D - None

Q 10 - The Index in Hive can be seen by

- A - Describe Index
- B - Show Index
- C - View Index
- D - Find Index

Q 11 - If an Index is dropped then

- A - The underlying table is also dropped
- B - The underlying table is not dropped
- C - the directory containing the index is deleted
- D - Error is thrown by hive

Q 12 - Indexes can be created

- A - only on managed tables
- B - only on views
- C - Only on external tables
- D - only on views with partitions

Q 13 - The clause " WITH DEFERRED REBUILD" while creating an index

- A - creates index on a table which is yet to be created
- B - creates index on a table which has no data
- C - creates index only on a table which has data
- D - creates an index which is empty

Q 14 - If the data on the table on which an index is defined changes then,

- A - The Index becomes invalid
- B - The index rebuilds automatically
- C - The Index has to be rebuilt manually
- D - The index must be dropped

Q 15 - The identifiers in HiveQL are

- A - case sensitive
- B - case insensitive
- C - sometimes case sensitive
- D - Depends on the Hadoop environment

Q 16 - What is the disadvantage of using too many partitions in Hive tables?

- A - It slows down the namenode
- B - Storage space is wasted
- C - Join queries become slow
- D - All of these

Q 17 - When importing data using SerDe, if a row is found to have more columns than expected then

- A - The extra columns are replaced with NULL
- B - The row is skipped
- C - The import halts with error
- D - The Columns are ignored.

Q 18 - Consider the below two sets of queries.

Query A:

```
hive> INSERT OVERWRITE TABLE sales
  SELECT * FROM history WHERE action = 'purchased';
hive> INSERT OVERWRITE TABLE credits
  SELECT * FROM history WHERE action = 'returned';
```

and

Query B:

```
hive> FROM history
  INSERT OVERWRITE sales SELECT * WHERE action = 'purchased'
  INSERT OVERWRITE credits SELECT * WHERE action = 'returned'
```

Which of them will make a single pass through?

- A - Query A
- B - Query B
- C - Both
- D - None

Q 19 - Which of the following feature is used to analyze the query execution plan

- A - Execute
- B - Planner
- C - Explain
- D - Optimizer

Q 20 - The LIMIT clause applied to a select query

- A - Executes the entire query before the LIMIT clause
- B - Hive fetches one row at a time until it gets the required limit.
- C - Creates a temporary file and stores the query result there
- D - Does a random sampling on the rows

Q 21 - The default limit to the number of rows returned by a query can be done using which of the following parameter?

- A - hive.limit.optimize.enable
- B - hive.limit.enable.optimize
- C - hive.optimize.enable.limit
- D - hive.enable.limit.optimize

Q 22 - The Property that decides what is the maximum number of files that can be sampled during the use of the LIMIT clause is

- A - hive.limit.optimize.file.max
- B - hive.limit.optimize.limit.file
- C - hive.limit.optimize.file.restrict
- D - hive.limit.optimize.limit.most

Q 23 - Which of the following hint is used to optimize the join queries

- A - `/* joinlasttablename */`
- B - `/* joinfirsttablename */`
- C - `/* streamtablename */`
- D - `/* cacheabletablename */`

Q 24 - Setting the local mode execution to true causes

- A - All tasks are executed on data available closest to the namenode
- B - All tasks are executed only on a single machine
- C - All the data files are cached on a datanode before query execution
- D - Random data is used for query execution

Q 25 - Hive can automatically decide to run local mode by setting which of the following parameters in hive-site.xml?

- A - hive.exec.mode.local.enable
- B - hive.exec.mode.cluster.disable
- C - hive.exec.mode.local.first

ANSWER SHEET

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