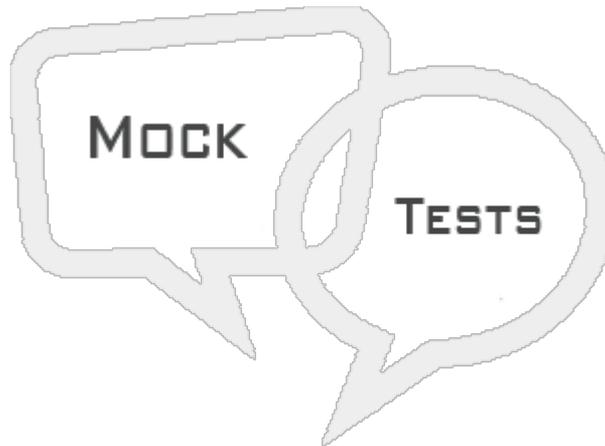


SQOOP MOCK TEST

<http://www.tutorialspoint.com>

Copyright © tutorialspoint.com

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



SQOOP MOCK TEST I

Q 1 - Which of the following is used by sqoop to establish a connection with enterprise data warehouses?

- A - RDBMS driver
- B - JDBC Driver
- C - IDBC Driver
- D - SQL Driver

Q 2 - Besides the JDBC driver, sqoop also needs which of the following to connect to remote databases?

- A - Putty
- B - SSH
- C - Conenctor
- D - sqoop client

Q 3 - To run sqoop from multiple nodes, it has to be installed in

- A - Any one of the in the local filesystem.
- B - each of the node where it is supposed to run
- C - Only on a pair of nodes of the cluster
- D - Need not be installed.

Q 4 - By default the records from databases imported to HDFS by sqoop are

- A - Tab separated

B - Concatenated columns

C - space separated

D - comma separated

Q 5 - To import data to Hadoop cluster from relational database sqoop create a mapreduce job. In this job

A - All the data is transferred in one go.

B - each mapper transfers a slice of Table's data

C - Each mapper transfers tables' data along with table's metadata *nameofthecolumnsetc*

D - Only the schema of relational table is validated without fetching data

Q 6 - The parameter in sqoop which specifies the output directories when importing data is

A - --output-path

B - --target-path

C - --output-dir

D - --target-dir

Q 7 - If there is already a target directory with the same name as the table being imported then

A - The directory gets deleted and recreated.

B - The sqoop job fails

C - Another directory under the existing directory gets created.

D - The existing directory gets renamed

Q 8 - To prevent the password from being mentioned in the sqoop import clause we can use the additional parameters

A - -p

B - --password-file

C - both of these

D - cannot be prevented

Q 9 - What are the two binary file formats supported by sqoop?

A - Avro & SequenceFile

B - Rcf file and SequenceFile

C - ORC file and RC file

D - Avro and RC file

Q 10 - While SequenceFile stores each record as key-value pair, the avro system stored records as

- A - Simple text
- B - chained lists
- C - Linked lists
- D - schema and data

Q 11 - The compression mechanism used by sqoop is

- A - built in sqoop
- B - delegated to Hadoop
- C - supplied as a java plugin to sqoop
- D - Needs to be installed in the OS running sqoop

Q 12 - For some databases sqoop can to faster data transfer by using the parameter

- A - --bulkload
- B - --fastload
- C - --dump
- D - --direct

Q 13 - The data type mapping between the database column and sqoop column can be overridden by using the parameter

- A - --override-column-type
- B - --map-column-type
- C - --override-column-java
- D - --map-column-java

Q 14 - What does the num-mappers parameter serves?

- A - force sqoop to use only one map task
- B - set the number of map tasks sqoop can use
- C - store the data imported by each map tasks in a separate file
- D - Fetch each row from the table using a new map task

Q 15 - What is the default value used by sqoop when it encounters a missing value while importing from CSV file.

A - NULL

B - null

C - space character

D - No values

Q 16 - What option can be used to import the entire database from a relational system using sqoop?

A - --import-all-db

B - --import-all-tables

C - --import-all

D - --import

Q 17 - what option can be used to import only some of the table from a database while using the --import-all-tables parameter?

A - --skip-tables

B - --without-tables

C - --forgo-tables

D - --exclude-tables

Q 18 - Sqoop supports

A - full import of tables

B - partial import of data from tables

C - Both full and partial data import

D - Import both the table and its partitions

Q 19 - What are the two different incremental modes of importing data into sqoop?

A - merge and add

B - append and modified

C - merge and lastmodified

D - append and lastmodified

Q 20 - What does the --last-value parameter in sqoop incremental import signify?

A - What is the number of rows successfully imported in append type import

B - what is the date value to be used to select the rows for import in the last_update_date type import

C - Both of the above

D - The count of the number of rows that were successful in the current import.

Q 21 - The --options-file parameter is used to

- A - save the import log
- B - specify the name of the data files to be created after import
- C - store all the sqoop variables
- D - store the parameters and their values in a file to be used by various sqoop commands.

Q 22 - while specifying the connect string in the sqoop import command, for a Hadoop cluster, if we specify localhost in place of a server address *hostnameorIPaddress* in the URI, then

- A - The import job will connect to local databases
- B - Each node may connect to different databases
- C - the import job may succeed
- D - All of the above

Q 23 - What is the disadvantage of storing password in the metastore as compared to storing in a password file?

- A - it is easily accessible
- B - it may get deleted accidentally
- C - It cannot be updated
- D - it is unencrypted

Q 24 - What is the advantage of storing password in a metastore as compared to storing in password in a file?

- A - It can be run by any user with valid access to sqoop environment
- B - The password in metastore can be updated while that in password file cannot be
- C - The password file can be encrypted while the metastore cannot be encrypted
- D - User intervention is required in password file but not in metastore.

Q 25 - The argument in a saved sqoop job can be altered at run time by using the option

- A - --alter
- B - --newval
- C - --exec
- D - --changeparam

Q 26 - What is achieved by using the --meta-connect parameter in a sqoop command?

- A - run metastore as a service accessible remotely
- B - run metastore as a service accessible locally
- C - connect to the metastore tables
- D - connect to the metadata of the external relational tables from which data has to be imported

Q 27 - The free-form query import feature in sqoop allows to import data from

- A - non relations sources
- B - a relational source without using a connector
- C - a relation source using a sql query
- D - a relational source using custom java classes

Q 28 - The clause 'WHERE \$CONDITIONS' in the sql query specified to import data, serves the purpose of

- A - split the query result into multiple chunks while importing
- B - picking a subset of rows from a table
- C - specify the database from where the data needs to be imported
- D - Specify the target directory where the data will be stored.

Q 29 - The parameter to give a custom name to the mapreduce job running a sqoop import command is –

- A - --sqoop-job-name
- B - --map-job-name
- C - --mapreduce-job-name
- D - --rename-job

Q 30 - While using a free-form query to import data, Sqoop finds that two columns from the joined tables have the same name. In this case the job

- A - will fail
- B - will run ignoring the column from each table
- C - will prompt the user to rename one of the columns
- D - automatically create an alias for one of the columns as succeed the job.

Q 31 - The -boundary-query parameter is used to

- A - Select the maximum number of rows to be retrieved by the query
- B - Select maximum and minimum values of the column specified in the -split-by parameter

C - Select the number of splits they query can run

D - Select the maximum and minimum number of mapreduce tasks that will be used in the query.

Q 32 - In a table import the name of the mapreduce job

A - Is named after the table name

B - Can be customized

C - Can be passed as a query parameter

D - Is a random name decided by the system.

Q 33 - In the import involving join of two tables the if there are two columns with matching name between two tables then this conflict can be resolved by

A - Using table aliases

B - Column aliases

C - First creating temporary tables form each table with different column names

D - Rename the columns in the source system and then import

Q 34 - Data Transfer using sqoop can be

A - only imported into the Hadoop system

B - both imported and exported from Hadoop system

C - transformed during import

D - transformed during the export

Q 35 - While importing data into Hadoop using sqoop the SQL SELCT clause is used. Similarly while exporting data form Hadoop the SQL clause used is

A - APPEND

B - MERGE

C - UPDTAE

D - INSERT

Q 36 - While inserting data into Relational system from Hadoop using sqoop, the various table constraints present in the relational table must be

A - Disabled temporarily

B - Dropped and re created

C - Renamed

D - Not violated

Q 37 - The export and import of data between sqoop and relational system happens through which of the following programs?

- A - Sqoop client program
- B - Mapreduce job submitted by the sqoop command
- C - Database stores procedure
- D - Hdfs file management program

Q 38 - When does sqoop gather the metadata of the relational table into which it exports the data?

- A - Gathers the metadata of all tables only once during establishing the connection to the database
- B - Never as it relies on the user to ensure the exported data matches the table's structure
- C - Every time the sqoop export command is submitted and just before the data transfer starts.
- D - Only if the export fails, Sqoop accesses the metadata of the table

Q 39 - Sqoop's default behavior while inserting rows into relational tables is

- A - one row at a time
- B - multiple rows depending on the memory availability
- C - It depends on the database driver being used
- D - Executes random number of insert statements depending on the CPU availability

Q 40 - Which parameter in sqoop is used for bulk data export to relational tables?

- A - -bulk
- B - -batch
- C - -load
- D - -grouped data

Q 41 - What does the parameter " Dsqoop.export.records.per.statemet=10" do in a sqoop export command?

- A - Exports 10 records in each insert statement
- B - Export 10 insert statements every time the command runs
- C - Export only the first 10 record to the table
- D - Run 10 data export commands in parallel

Q 42 - The parameter which decided How many rows will be inserted per transaction in sqoop is

- A - Dsqoop.export.rows.per.transaction

- B - Dsqoop.export.records.per.transaction
- C - Dsqoop.export.inserts.per.transaction
- D - Dsqoop.export.statements.per.transaction

Q 43 - The insert query used to insert exported data into tables is generate by

- A - Sqoop command and processed as such
- B - Sqoop command and modified suitably by JDBC drive
- C - JDBC driver
- D - Database specific driver

Q 44 - When the “sqoop.export.records.per.statement” is set to two or more, the query created by sqoop has the SQL forma of

- A - INSERT INTO TABLE VLAUES. . ;NSERT INTO TABLE VLAUES. . ; and so on.
- B - BULK NSERT INTO TABLE VLAUES. . ,,,,,,;
- C - NSERT INTO TABLE VLAUES. . ,VALUES. . ,VALUES.....
- D - NSERT INTO TABLE VLAUES. . ,,,,,,;

Q 45 - What happens if the sqoop generated export query is not accepted by the database?

- A - The export fails
- B - The export succeeds partially.
- C - The export does not start
- D - Sqoop automatically modifies the query to succeed after receiving the failure response form database.

Q 46 - Using the higher value for the parameter sqoop.export.statements.per.transaction will

- A - Always increase the export performance
- B - May or may not increase the export performance
- C - May decrease the performance for some of the tables
- D - Cause frequent commits to the database

Q 47 - The -staging-table parameter is used for

- A - Storing some sample data from Hadoop before loading the real table
- B - Storing all the required data from Hadoop before loading it to real table
- C - Storing the rejected rows

D - Storing the metadata structure of tables to which data is being exported

Q 48 - With the -staging-table parameter, the data is moved from staging to final table

- A - Automatically if staging load is successful
- B - Has to be done by user after verifying the data in staging
- C - Depends on the data size
- D - Depends on the memory available to move the data

Q 49 - Which of the following is a disadvantage of using the -staging-table parameter?

- A - Data is stored twice and consumes more memory
- B - The overall export time is more than direct export to final table
- C - User should ensure the structure of staging table and final tables are in Sync.
- D - All of the above

Q 50 - Using the -staging-table parameter while loading data to relational tables the creation of staging table is done

- A - Automatically by sqoop
- B - Automatically by database
- C - User has to ensure it is created
- D - Automatically created by a Hadoop process beyond sqoop

ANSWER SHEET

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

- | | |
|----|---|
| 1 | B |
| 2 | C |
| 3 | A |
| 4 | D |
| 5 | B |
| 6 | D |
| 7 | B |
| 8 | C |
| 9 | A |
| 10 | D |
| 11 | B |

12	D
13	D
14	B
15	B
16	C
17	D
18	C
19	D
20	C
21	D
22	D
23	D
24	A
25	C
26	A
27	C
28	A
29	C
30	A
31	B
32	A
33	B
34	B
35	D
36	D
37	B
38	C
39	A
40	B
41	A
42	D
43	A
44	D
45	A
46	C

47	B
48	A
49	D
50	C

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