SQOOP MOCK TEST

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.

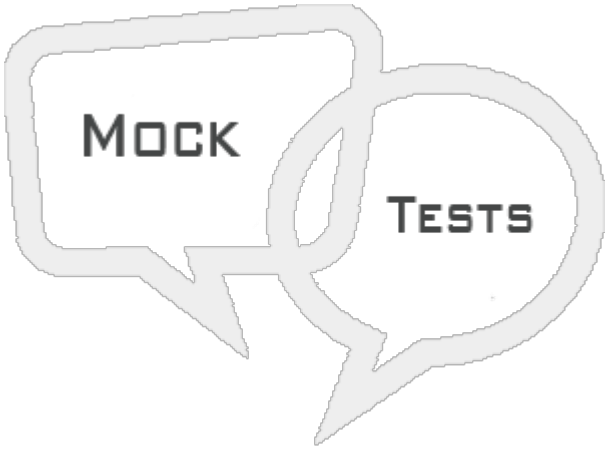

SQOOP MOCK TEST III

Q 1 - Sqoop can automatically clear the staging table before loading by using the parameter

A - –clear-table

B - –clear-staging-table

C - --truncate-staging-table

D - -delete-from-staging-table

Q 2 - Can sqoop use the TRUNCATE option in database while clearing data from a table?

A - Yes

B - No

C - Depends on the database

D - Depends on the Hadoop configuration

Q 3 - The –update-key parameter is used to

A - Update the primary key field present in the Hadoop data to be exported

B - Update the primary key field in the table to which data is already exported

C - Update the database connectivity parameters like username, password etc

D - Update the already exported rows based on a primary key field

Q 4 - The –update-key parameter can take

A - Only one column name as key field

B - Two column name as key filed

C - Comma separated list of columns as keys

D - Table name and column name as key

Q 5 - A table contains 4 columns $C1, C2, C3, C4$. With –update-key C2,C4, the sqoop generated query will be like

A - Update table set C1 = ‘newval’, c3 = ’newval’ where c2 = ‘oldval’ and c4 = ’oldval’

B - Update table set C2 = ‘newval’, c4 = ’newval’ where c2 = ‘oldval’ and c4 = ’oldval’

C - Update table set C1 = ‘newval’, c2 = ’newval’, c3 = ’newval’, c4 = ’newval’ where c2 = ‘oldval’ and c4 = ’oldval’

D - None

Q 6 - The –update-key parameter can

A - Not insert new rows to the already exported table

B - Insert new rows to an already exported table

C - Insert new rows into the exported table only if it has a primary key

D - Delete rows form already exported table

Q 7 - Sqoop can insert new rows and update existing changed rows into an already exported table by using the parameter

A - –update-insert

B - -update-else-insert

C - –update-mode insert

D - –update-mode allowinsert

Q 8 - When using –update-mode allowinsert parameter with oracle database the feature of oracle used by sqoop is

A - UPSERT statement

B - MERGE statement

C - MULTITABLE INSERT statement

D - BULK LOAD statement

Q 9 - With MySQL, the feature used by sqoop for update or insert data into an exported table is

A - ON DUPLICATE KEY UPDATE

B - ON KEY UPDATE

C - ON NEW KEY UPDATE

D - ON NEW UPDATE

Q 10 - Can the upsert feature of sqoop delete some data form the exported table?

A - Yes

B - No

C - Depends on database

D - Yes With some additional sqoop parameters

Q 11 - To sync a HDFS file with some deleted rows with a previously exported table for the same table the option is to

A - Using staging table

B - Export the data again to a new database table and rename it

C - Use a ETL tool

D - Can not be done using sqoop

Q 12 - The parameter which can be used in place of --table parameter to insert data into table is

A - –call

B - –insert-into

C - –populate

D - –load-into

Q 13 - The disadvantage of using a stored procedure to laod data is

A - Data gets loaded faster

B - Parallel loads in the database table

C - The store procedure cannot load multiple tables at a time

D - It induces heavy load on database

Q 14 - If the table to which data is being exported has more columns than the data present in the hdfs file then

A - The load definitely fails

B - The load can be done only for the relevant columns present in HDFS file

C - The will populate values into the wrong columns

D - The load does not start

Q 15 - The parameter to specify only a selected number of columns to be exported to

a table is

A - -columns

B - –column-subset

C - ----columns-not-all

D - -columns-part

Q 16 - Load all or load nothing semantics is implemented by using the parameter

A - -loadd-all-nothing

B - -stage-load

C - -all-load

D - -staging-table

Q 17 - How do we decide the order of columns in which data is loaded to the target table?

A - By using -- order by parameter

B - By using a new mapreduce job aftet submitting sqoop export command

C - By using a database stored procedure

D - By using –columns parameter with comma separated column names in the required order.

Q 18 - What is the disadvantage of using the –columns parameter to insert a subset of columns to the relational table?

A - The relational table may have not null columns not covered in the –columns parameter.

B - The relational table may store the data from HDFS in wrong columns.

C - It may not load all the required data

D - It will not be able to populate primary key values

Q 19 - The parameter used to override NULL values to be inserted into relational targets is

A - -override-null

B - –input-null-string

C - -substitute-null

D - --replace-null

Q 20 - For Text based columns the parameter used for substituting null values is

A - -input-null-string

B - -input-null-non-string

C - -input-null-text

D - -input-null-varchar

Q 21 - For a column of data type numeric, the parameter used for substituting null values is

A - -input-null-string

B - -input-null-non-string

C - -input-null-text

D - -input-null-varchar

Q 22 - When a column value has a different data type in the HDFS system than expected in the relational table to which data will be exported −

A - Sqoop skips the rows

B - Sqoop fails the job

C - Sqoop loads the remaining rows by halting and asking whether to continue the load

D - Sqoop automatically changes the data type to a compatible data type and loads the data.

Q 23 - The parameter used in sqoop to import data directly into hive is

A - -import-direct

B - -import-hive

C - –hive-import

D - -hive-sqoop

Q 24 - While importing directly to hive using sqoop, if the table meta data does not exist in hive then

A - the sqoop job fails

B - sqoop creates the meta data in hive

C - sqoop waits for user to input the meta data

D - sqoop imports the data as a file without creating any meta data

Q 25 - To ensure that the columns created in hive by sqoop have the correct data types the parameter used by sqoop is

A - --map-column-hive

B - --map-column

C - --column-hive

D - --map-table-hive

Q 26 - During import to hive using sqoop the data is

A - directly loaded to existing hive table

B - first moved into a hive directory as a hdfs file

C - first moved into any temporary location as a HDFS file

D - First validated against the table schema in hive

Q 27 - While importing data to hive using sqoop, if data already exists in hive table then the default behaviour is

A - The incoming data is appended to hive table

B - the incoming data replaces data in hive table

C - The data only gets updated using the primary key of the hive table

D - sqoop command fails

Q 28 - To overwrite data present in hive table while importing data using sqoop, the sqoop parameter is

A - --sqoop-overwrite

B - --hive-overwrite

C - --hive-delete-all

D - --sqoop-delete-all

Q 29 - The temporary location to which sqoop moves the data before loading into hive is specified by the parameter

A - --target-dir

B - --source-dir

C - --hive-dir

D - --sqoop-dir

Q 30 - If the target hive table is partitioned then sqoop behavior is which of the following?

A - not load data into hive partitions

B - jobs will always fail

C - sqoop command will halt for user input for partition names

D - load data into hive partitions by using additional parameters

Q 31 - The parameter*s* used to laod data using sqoop into the hive partitions is/are

A - --hive-partition-key and -hive-partition-value

B - --hive-partition-key-value

C - --hive-partition-sequence

D - --hive-sqoop-partition-value

Q 32 - The data type of the column used for partition name while importing data using sqoop ino hive can be

A - only string

B - string or numeric

C - only date

D - string without special characters

Q 33 - The parameter --hive-drop-import-delims does which of the following?

A - replaces the hive delimiters with sqoop delimiters

B - drops the rows which do not have the \n,\t,\01 delimiters

C - removes all the \n,\t and \01 characters

D - drops the columns which do not have the \n,\t,\01 delimiters

Q 34 - The purpose of --hive-delims-replacement parameter in sqoop is to

A - Replace any hive delimiters with special string

B - Replace all the hive delimiters with null

C - replace \n, \t, and \01 characters with any other string

D - replace \n, \t, and \01 characters with null

Q 35 - HIve shows more row count than imported by sqoop. What can be the reason?

A - the \n chara cter present int the data

B - Error with java classes used in sqoop

C - limitation of features in hive

D - Insufficient memory issue

Q 36 - The parameter --hive-import can be used with

A - only importing to hive

B - importing to hive as well as text file

C - importing to hive and relational tables

D - importing to hive and binary data

Q 37 - To import data to HBase using sqoop the parameter*s* required is/are

A - --hbase-table

B - --hbase-columns

C - --hbase-table and --column-family

D - --habse-table and --hbase-column-family

Q 38 - If the hbase table to which sqoop is importing data does not exist then

A - sqoop creates it

B - sqoop job fails

C - sqoop waits for user input for hbase table details to proceed with import

D - sqoop imports the data to a temporary location under Hbase

Q 39 - The parameter used to identify the individual row in HBase while importing data to it using sqoop is

A - --hbase-row-key

B - --hbase-rowkey

C - --hbase-rowid

D - --hbase-row-id

Q 40 - The parameter that can create a hbase table using sqoop when importing data to hbase is

A - -hbase-create-table

B - -create-hbase-table

C - -create-hbase-table-columnlist

D - -create-hbase-table-rowkey

Q 41 - After importing a table into HBAse you find that the number of rows inserted is fewer than in the source. The possible reason is −

A - Sqoop is yet to have mature code for HBase

B - Sqoop version and Hbase version conflict

C - Hbase does not allow rows will all NULL values to be inserted

D - Hbase has very limited capabilities to handle numeric data types so some rows got rejected.

Q 42 - The property in sqoop that allows rows with all NULL values to be inserted into HBAse tables is −

A - sqoop.hbase.add.row.key

B - sqoop.hbase.allow.row.nulls,

C - Sqoop.hbase.add.row.id

D - It is not possible as HBAse will never allow rows with all null Columns to be inserted

Q 43 - When inserting data using sqoop into Hbase table in one physical node, the different parallel tasks of sqoop import create a bottleneck. This can be solved by

A - Configuring sqoop not to run parallel tasks

B - Configuring gHBase to accept rows in parallel

C - Creating more regions in HBAse table

D - Cannot be solved

Q 44 - The parameters in sqoop command can be passed in to Oozie by using which tags?

A - <parameters>

B - <args>

C - <sqoop>

D - <command>

Q 45 - In both import and export scenario, the role of ValidationThreshold is to determine if

A - the error margin between the source and target is within a range

B - the Sqoop command can handle the entire number of rows

C - the number of rows rejected by sqoop while reading the data

D - the number of rows rejected by the target database while loading the data

Q 46 - The comparison of row counts between the source system and the target database while loading the data using sqoop is done using the parameter

A - --Validate

B - --Rowcount

C - -row*count*

D - -allrows

Q 47 - The sqoop export/import jobs canbe stored and used again and again by using

A - sqoop- jobname

B - sqoop-save-job

C - sqoop-all-jobs

D - sqoop-job

Q 48 - What is achieved by the command – sqoop job –exec myjob

A - Sqoop job named myjob is saved to sqoop metastore

B - Sqoop job named myjob starts running

C - Sqoop job named myjob is scheduled

D - Sqoop job named myjob gets created

Q 49 - The tool in sqoop which combines two data sets and preserves only the latest values using a primary key is

A - sqoop-merge

B - sqoop-combine

C - sqoop-preserve

D - sqoop-distinct

Q 50 - The tool that populates a Hive metastore with a definition for a table based on a database table previously imported to HDFS is

A - create-hive-table

B - import-hive-metastore

C - create-hive-metastore

D - update-hive-metastore

ANSWER SHEET

Question Number	Answer Key
1	B
2	C
3	D
4	C
5	A
6	A
7	D
8	B
9	A
10	A
11	B
12	A

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

48	B
49	A
50	B