CCB-400 Q&As

Cloudera Certified Specialist in Apache HBase

Pass Cloudera CCB-400 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/ccb-400.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Cloudera Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

You have two tables in an existing RDBMS. One table contains order information (item, quantity, price, etc.) and the other contains store information (address, phone, manager, etc). These two tables are not often accessed simultaneously. You would like to move this data into HBase. How would you design the schema?

A. Create two tables each with a single columnfamily

B. Create a single table with one column family

C. Create a single table with two column families

D. Create two tables each with multiple column families

Correct Answer: C

QUESTION 2

Your client is writing to a region when the RegionServer crashes. At what point in the write is your data secure?

A. From the moment the RegionServer wrote to the WAL (write-ahead log)

B. From the moment the RegionServer returned the call

C. From the moment the RegionServer received the call

D. From the moment the RegionServer wrote to the MemStore

Correct Answer: A

QUESTION 3

The cells in a given row have versions that range from 1000 to 2000. You execute a delete specifying the value 3000 for the version. What is the outcome?

A. The delete fails with an error.

B. Only cells equal to the Specified version are deleted.

C. The entire row is deleted.

D. Nothing in the row is deleted.

Correct Answer: C

Reference:http://archive.cloudera.com/cdh4/cdh/4/hbase/book.html#delete(scroll below and see 5.8.1.5. Delete topic, read the last paragraph)

QUESTION 4

You have a table with 5 TB of data, 10 RegionServers, and a region size of 256MB. You want to continue with puts to widely disbursed row ids in your table. Which of the following will improve write performance?

A. Increase your buffer cachein theRegionServers

B. Increase the number of RegionServers to 15

C. Decrease your number of RegionServersto 5

D. Decreaseyour regionsize to 128MB

Correct Answer: C

QUESTION 5

You want to do mostly full table scans on your data. In order to improve performance you increase your block size. Why does this improve your scan performance?

A. It does not. Increasing block size does not improve scan performance.

B. It does not. Increasing block size means that fewer blocks fit into your block cache. This requires HBase to read each block from disk rather than cache for each scan, thereby decreasing scan performance.

C. Increasing block size requires HBase to readfrom disk fewer times, thereby increasing scan performance.

D. Increasing block size means fewer block indexes that need to be read from disk, therebyincreasing scan performance.

Correct Answer: D

QUESTION 6

You have a "Users" table in HBase and you would like to insert a row that consists of a UserID, "jsmith70" and an email address, "jane@example.com". The table has a single Column Family named "Meta" and the row key will be the user\\'s ID. The shell command you should use to complete this is:

A. put`Users\\', `jsmith70\\',`jane@example.com\\'

B. put `Users\\', `UserID:jsmith70\\', `Email:jane@example.com\\'

C. put `Users\\', `jsmith70\\', `Meta:Email\\', `jane@example.com\\'

D. put `Users\\',`Meta:UserID\\',`jsmith70\\',`Meta:Email,`jane@example.com\\'

Correct Answer: B

QUESTION 7

You have a total of three tables stored in HBase. Exchanging catalog regions, how many regions will your RegionServers have?

A. Exactly three

B. Exactly one

C. At least one

D. Atleastthree

Correct Answer: B

QUESTION 8

Given that following is your entire dataset:

100 column=Managers:Name, timestamp=13313141762084, value=Steve 100 column=Manage:Salary, timestamp=13313141762086, value=80000 100 column=Skills:Skill_1, timestamp=13313141762089, value=Hadoop 100 column=Skills:Skill_2, timestamp=13313141762092, value=HBase How many sets of physical files will be read during a scan of the entire dataset immediately following a major compaction?

A. Two

B. One

C. Three

D. Four

Correct Answer: B

QUESTION 9

Your client connects to HBase for the first time to road a row user_1234 located in a table Users. What process does your client use to find the correct RegionServer to which it should send the request?

A. The client looks up the location of ROOT,in which it looks up the location of META, in whichit looks up the location of the correct Users region.

B. The client looks up the location of the master, in which it looks up the location of META,in which it looks up thelocation of the correct Users region.

C. The client looks up the location of ROOT in which it looks up the location of the correctUsersregion.

D. The client queries the master to find the location of the Users table.

Correct Answer: B

QUESTION 10

Given the following HBase table schema:

Row Key, colFam_A:a, colFam_A:b, colFamB:2, colFam_B:10

A table scan will return the column data in which of the following sorted orders:

A. Row Key, colFam_A:a, colFam__A:b, colFam_B:10, colFam_B:2

B. Row Key, colFam_A:a, colFam__A:b,colFam_B:2, colFam_B:10

C. Row Key, colFam_A:a, colFam__B:2, colFam_A:b, colFam_B:10

D. Row Key, colFam_A:a, colFam__B:10, colFam_A:b, colFam_B:2

Correct Answer: B