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.pass2lead.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 need to free up disk space on your HBase cluster. You delete all versions of your data that is older than one week. You notice your delete has had minimal impact on your storage availability. This is because:

A. You have large store file indexes

B. HBase has not flushed the MemStore

C. HBasehas not run a minor compaction

D. HBase has notrun a major compaction

Correct Answer: A

QUESTION 2

Data is written to the HLog in which of the following orders?

A. In order of writes

B. In order of writes, separated by region

C. Ascending first by region and second by row key

D. Descending first by region and second by row key

Correct Answer: D

QUESTION 3

You have a table where keys range from "A" to "Z", and you want to scan from "D" to "H." Which of the following is true?

A. A MultiGet must be issued for rows D, E, F, G, H.

B. The scan class supports ranges via the stop and start rows.

C. All scans are full table scans, the client must implement filtering.

D. In order to range scan, raw scan mode must be enabled.

Correct Answer: B

QUESTION 4

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 5

You have one primary HMaster and one standby. Your primary HMaster Falls fails and your client application needs to make a metadata change. Which of the following is the effect on your client application?

A. The client will queryZooKeepertofind the location of the new HMaster and complete the metadata change.

B. The client will make the metadata change regardless of the slate of the HMaster.

C. The new HMaster will notify the client and complete the metadata change.

D. Theclientapplication will fail with a runtime error.

Correct Answer: A

QUESTION 6

Which of the following configuration values determines automated splitting?

A. hbase.hregion.majorcompaction

B. hbase.hregion.flush.size

C. hbase.balancer.period

D. hbase.hregion.max.filesize

Correct Answer: D

Reference:http://hbase.apache.org/book/important_configurations.html(2.8.2.6. Bigger regions, see the code in the last sentence)

QUESTION 7

Given that the 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 regions will be read during a scan of the entire dataset?

A. Four

B. Two

C. One

D. Three

Correct Answer: A

QUESTION 8

You want to do a full table scan on your data. You decide to disable block caching to see if this improves scan performance. Will disabling block caching improve scan performance?

A. No. Disabling block caching does notimprove scan performance.

B. Yes.When you disable blockcaching, youfreeup that memory for other operations.Witha full table scan,you cannot take advantage of blockcaching anyway because your entiretable won\\'tfitinto cache.

C. No. If you disable block caching, HBase must read each block index from disk for each scan,thereby decreasing scan performance.

D. Yes. When you disable block caching, youfree up memory forMemStore, which improves,scanperformance.

Correct Answer: B

QUESTION 9

Your client application needs to scan s region for the row key value 104.

Given a store that contains the following list of Row Key values:

100, 101, 102, 103, 104, 105, 106, 107

A bloom filter would return which of the following?

A. Confirmation that 104 may becontained in theset

B. Confirmation that 104 is contained in the set

C. Thehash of column family

D. The file offset of the value 104

Correct Answer: B

QUESTION 10

Your client application if; writing data to a Region. By default, where is the data saved first?

A. StoreFile

B. WAL

C. MemStore

D. Local disk on theRegionServer

Correct Answer: C

Reference:http://www.cloudera.com/blog/2012/07/hbase-log-splitting/(Log splitting, first paragraph)