Vendor: Cloudera

Exam Code: CCB-400

Exam Name: Cloudera Certified Specialist in Apache HBase

Version: Demo

QUESTION NO: 1

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.

Answer: A
Explanation:

QUESTION NO: 2

You have an average key-value pair size of 100 bytes. Your primary access is random needs on the table. Which of the following actions will speed up random reading performance on your cluster?

A. Turn off WAL on puts
B. Increase the number of versions kept
C. Decrease the block size
D. Increase the block size

Answer: D
Explanation:

QUESTION NO: 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.

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 NO: 4

You have an "Employees" table in HBase. The Row Keys are the employees' IDs. You would like to retrieve all employees who have an employee ID between 'user_100' and 'user_110'. The shell command you would use to complete this is:

A. scan'Employees', {STARTROW =>'user_100', STOPROW =>'user_111'}
B. get'Employees', {STARTROW =>'user_100', STOPROW =>'user_110'}
C. scan'Employees', {STARTROW =>'user_100', SLIMIT => 10}
D. scan'Employees', {STARTROW =>'user_100', STOPROW =>'user_110'}

Answer: D
Explanation:

QUESTION NO: 5

Under default settings, which feature of HBase ensures that data won't be lost in the event of a RegionServer failure?

A. AllHBase activity is written to the WAL, which is stored in HDFS
B. All operations are logged on theHMaster.
C. HBase is ACID compliant, which guarantees that itis Durable.
D. Data is stored on the local filesystem of the RegionServer.

Answer: A
Reference:http://tm.durusau.net/?p=27674(See 'From the post' second paragraph)

QUESTION NO: 6

You have two standbys and one primary HMaster. Your primary HMaster fails. Which of the remaining HMasters becomes the new primary?

A. Whichever HMaster first responds to ZooKeeper
B. Whichever HMasterZooKeeper randomly selects
C. Whichever HMastercreates the znode first
D. Whichever HMaster has the lower IP address

Answer: A
Explanation:

QUESTION NO: 7

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

Answer: D
Explanation:

QUESTION NO: 8

You have a table with the following rowkeys:

r1, r2, r3, r10, r15, r20, r25, r30, r35

In which order will these rows be retrieved from a scan?

A. r35,r30,r3,r25,r20,r2,r15,r10,r1
B. r1,r2,r3,r10,r15,r20, r25,r30,r35
C. r1,r10,r15,r2,r20,r25,r3,r30,r35
D. r35,r30,r25,r20,r15,r10,r3,r2,r1

Answer: D
Explanation:

QUESTION NO: 9

You need to create a "WebLogs" table in HBase. The table will consist of a single Column Family called "Errors" and two column qualifiers, "IP" and "URL". The shell command you should use to create the table is:

A. create'WebLogs',{NAME =>'Errors:IP', NAME =>'Errors:URL'}
B. create'WebLogs','Errors'{NAME =>'IP', NAME =>'URL'}
C. create'WebLogs','Errors:IP','Errors:URL'
D. create'WebLogs','Errors'

Answer: C
Explanation:

QUESTION NO: 10

Which feature of HBase ensures predictable disk head seek performance within a RegionServer?

A. Data is stored distributed in HDFS
B. Data stored in HBase is sparse
C. Data is stored sorted on row keys
D. Data is stored as an uninterpreted array of bytes

Answer: C
Explanation:

QUESTION NO: 11

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

Answer: B
Explanation:

QUESTION NO: 12

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

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

QUESTION NO: 13

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.

Answer: B

Explanation:

QUESTION NO: 14

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

Answer: B
Explanation:

QUESTION NO: 15

You have 40 Web servers producing timeseries data from Web traffic logs. You want to attain high write throughput for storing this data in an HBase table. Which of these should you choose for a row key to maximize your write throughput?

A. <hashCode(centralServerGeneratedSequenceID) ><timestamp>
B. <Long.MAX_VALUE – timestamp>
C. <timestamp>
D. <hashCode(serverGeneratingTheWeblog)><timestamp>

Answer: C
Explanation:

QUESTION NO: 16

You have data already stored in HDFS and are considering using HBase. Which additional feature

does HBase provide to HDFS?

A. Random writes
B. Batchprocessing
C. Fault tolerance
D. Scalability

Answer: A
Reference:http://borthakur.com/ftp/SIGMODRealtimeHadoopPresentation.pdf(11th slide)

QUESTION NO: 17

Given the following 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 store files will be contained in your region(s) immediately following a major compaction?

A. Four
B. Three
C. Two
D. One

Answer: D
Explanation:

QUESTION NO: 18

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 cloedete 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

Answer: A
Explanation:

QUESTION NO: 19

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

Answer: B
Explanation:

QUESTION NO: 20

Your HBase cluster has hit a performance wall and doesn't seem to be getting faster as you add

RegionServers. Adding an additional HMaster will:

A. Have no effect on performance.
B. Improve the performance of region writes but decreasethe performance of metadata changes.
C. Improve the performance of metadata chancier, but decrease the performance of region writes.
D. Make theperformance problem even worse, as operations will have to be replicated to multiplemasters.

Answer: A
Explanation:

To Read the **Whole Q&As**, please purchase the **Complete Version** from **Our website**.

Trying our product !

★ **100%** Guaranteed Success

★ **100%** Money Back Guarantee

★ **365 Days** Free Update

★ **Instant Download** After Purchase

★ **24x7** Customer Support

★ Average **99.9%** Success Rate

★ More than **69,000** Satisfied Customers Worldwide

★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:

One Year Free Update	Money Back Guarantee	Security & Privacy
Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.	To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.	We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.

Guarantee & Policy | Privacy & Policy | Terms & Conditions

Any charges made through this site will appear as Global Simulators Limited.
All trademarks are the property of their respective owners.