



# CCB-400<sup>Q&As</sup>

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.geekcert.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 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\' }`

Correct Answer: D

---

### QUESTION 2

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

Correct Answer: C

---

### QUESTION 3

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 4

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 5

Your data load application is maintaining a custom versioning scheme (not using the timestamp as the version number). You accidentally executed three writes to a given cell all with the same version during which time no flushes have occurred. Which of the three data writes will dBase maintain?

- A. None of the writes to cell
- B. The last write to cell
- C. The first write to cell
- D. All of the writes to cell

Correct Answer: C

[CCB-400 PDF Dumps](#)

[CCB-400 VCE Dumps](#)

[CCB-400 Braindumps](#)