

Unit 8 - Week 6

Course outline

How does an NPTEL online course work?

Week 0 Assignment 0

Week 1

Week 2

Week 3

Week 4

Week 5

Week 6

Design of Key-Value Stores

Design of HBase

Week 6 : Lecture Material

Quiz : Assignment-6

Week 6 - FEEDBACK

Assignment-6 Solution

Week 7

Week 8

Text Transcripts

DOWNLOAD VIDEOS

LIVE Session

Assignment-6

The due date for submitting this assignment has passed.
As per our records you have not submitted this assignment.

Due on 2020-03-11, 23:59 IST.

1) Apache Cassandra is a massively scalable open source _____ database. 1 point

- SQL
 NoSQL
 NewSQL
 None of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
NoSQL

2) Consider the following statements related to the replication strategies in Cassandra: 1 point

Statement 1: SimpleStrategy is highly recommended for most deployments because it is much easier to expand to multiple data centers when required by future expansion, it specifies how many replicas you want in each data center.

Statement 2: NetworkTopologyStrategy is used for a single data center only. It uses the Partitioner, of which there are two kinds: (i) RandomPartitioner for Chord-like hash partitioning (ii) ByteOrderedPartitioner to assign ranges of keys to servers.

- Only statement 1 is true
 Only statement 2 is true
 Both statements are true
 Both statements are false

No, the answer is incorrect.
Score: 0

Accepted Answers:
Both statements are false

3) _____ a memory cache to store the in memory copy of the data. It accumulates writes and provides read for data which are not yet stored to disk. 1 point

- Distributed Hash tables (DHT)
 Memtable
 Collection
 SSTable

No, the answer is incorrect.
Score: 0

Accepted Answers:
Memtable

4) True or False ? 1 point

Cassandra extends the concept of eventual consistency. Eventual consistency implies the storage system guarantees that if no new updates are made to the object, eventually all accesses will return the last updated value.

- True
 False

No, the answer is incorrect.
Score: 0

Accepted Answers:
True

5) Consider the following statements: 1 point

Statement 1: Scale up means grow your cluster capacity by replacing with more powerful machines .

Statement 2: Scale out means incrementally grow your cluster capacity by adding more COTS machines (Components Off the Shelf).

- Only statement 1 is true
 Only statement 2 is true
 Both statements are true
 Both statements are false

No, the answer is incorrect.
Score: 0

Accepted Answers:
Both statements are true

6) HBase is a distributed _____ database built on top of the Hadoop file system. 1 point

- Row-oriented
 Column-oriented
 Tuple-oriented
 None of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
Column-oriented

7) A small chunk of data residing in one machine which is part of a cluster of machines holding one Hbase table is known as _____ 1 point

- Rowarea
 Tablearea
 Region
 Split

No, the answer is incorrect.
Score: 0

Accepted Answers:
Region

8) _____ is a Cassandra feature that optimizes the cluster consistency process. 1 point

- Hinted handon
 Hinted handoff
 Tombstone
 Hinted tomb

No, the answer is incorrect.
Score: 0

Accepted Answers:
Hinted handoff

9) Cell in HBase Table is a combination of _____ 1 point

- Row and column family
 Row, column family and column qualifier
 Row, column family, column qualifier and row keys
 Row, column family, column qualifier and contains a value and a timestamp

No, the answer is incorrect.
Score: 0

Accepted Answers:
Row, column family, column qualifier and contains a value and a timestamp

10) In Cassandra, _____ job is to determine which data centers and racks it should use to read data from and write data to. 1 point

- Client requests
 Snitch
 Partitioner
 None of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
Snitch

11) Cassandra uses a protocol called _____ to discover location and state information 1 point

- Gossip
 HBase
 Key-value
 None of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
Gossip

12) Fill the correct choices for the given scenarios: 1 point

_____ Reads/writes complete reliably and quickly.

_____ When thousands of customers are looking to book a flight, all updates from any client (e.g.- book a flight) should be accessible by other clients.

_____ Can happen across datacenters when the Internet gets disconnected

- Availability, Consistency, Partition tolerance
 Consistency, Availability, Partition tolerance
 Partition tolerance, Consistency, Availability
 Consistency, Partition tolerance, Availability

No, the answer is incorrect.
Score: 0

Accepted Answers:
Availability, Consistency, Partition tolerance