

Unit 10 - Week 8

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

Week 7

Week 8

Introduction to Spark

Introduction to Kafka

Week 8: Lecture Material

Quiz : Assignment-8

Week 8 - FEEDBACK

Assignment-8 Solution

Text Transcripts

DOWNLOAD VIDEOS

LIVE Session

Assignment-8

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

Due on 2020-03-25, 23:59 IST.

1) Which is not a component on the top of Spark Core? 1 point

- Spark RDD
- Spark Streaming
- MLlib
- None of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
Spark RDD

2) True or False ? 1 point

RDDs (Resilient Distributed Datasets) are fault-tolerant and immutable.

- True
- False

No, the answer is incorrect.
Score: 0

Accepted Answers:
True

3) In Kafka, _____ is the node responsible for all reads and writes for the given partition. 1 point

- Leader
- Follower
- Stater
- None of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
Leader

4) Which type of processing Apache Spark can handle ? 1 point

- Batch Processing
- Stream Processing
- Graph Processing
- All of the Mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
All of the Mentioned

5) In Spark, a _____ is a read-only collection of objects partitioned across a set of machines that can be rebuilt if a partition is lost. 1 point

- Spark Streaming
- Resilient Distributed Dataset (RDD)
- FlatMap
- Driver

No, the answer is incorrect.
Score: 0

Accepted Answers:
Resilient Distributed Dataset (RDD)

6) Identify the correct kafka commands for the following: 1 point

P: It is the distributed, durable equivalent of Unix pipes. Use it to connect and compose your large-scale data applications.

Q: These are the commands of your Unix pipelines. Use it to transform data stored in Kafka.

R: It is the I/O redirection in your Unix pipelines. Use it to get your data into and out of Kafka.

- P: Kafka Streams, Q: Kafka Connect, R: Kafka Core
- P: Kafka Core, Q: Kafka Connect, R: Kafka Stream
- P: Kafka Streams, Q: Kafka Core, R: Kafka Connect
- P: Kafka Core, Q: Kafka Streams, R: Kafka Connect

No, the answer is incorrect.
Score: 0

Accepted Answers:
P: Kafka Core, Q: Kafka Streams, R: Kafka Connect

7) Kafka is run as a cluster comprised of one or more servers each of which is called _____. 1 point

- cTakes
- Broker
- Chunks
- None of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
Broker

8) Kafka maintains feeds of messages in categories called _____. 1 point

- Topics
- Chunks
- Domains
- Messages

No, the answer is incorrect.
Score: 0

Accepted Answers:
Topics

9) _____ is a distributed graph processing framework on top of Spark. 1 point

- MLlib
- Spark Streaming
- GraphX
- All of the mentioned

No, the answer is incorrect.
Score: 0

Accepted Answers:
GraphX

10) _____ is a distributed machine learning framework on top of Spark. 1 point

- MLlib
- Spark Streaming
- GraphX
- RDDs

No, the answer is incorrect.
Score: 0

Accepted Answers:
MLlib