

KillTest

High Quality. Better Service



Practice Questions

<https://www.killtest.com>

Free updates for one year

Exam : BDS-C00

**Title : AWS Certified BigData -
Specialty**

Version : DEMO

1.Using only AWS services. You intend to automatically scale a fleet of stateless of stateless web servers based on CPU and network utilization metrics.

Which of the following services are needed? Choose 2 answers

- A. Auto Scaling
- B. Amazon Simple Notification Service
- C. AWS Cloud Formation
- D. CloudWatch
- E. Amazon Simple Workflow Service

Answer: A,D

2.A company has several teams of analytics. Each team of analysts has their own cluster. The teams need to run SQL queries using Hive, Spark-SQL and Presto with Amazon EMR. The company needs to enable a centralized metadata layer to expose the Amazon S3 objects as tables to the analysts.

Which approach meets the requirement for a centralized metadata layer?

- A. EMRFS consistent view with a common Amazon DynamoDB table
- B. Bootstrap action to change the Hive Metastore to an Amazon RDS database
- C. s3distcp with the output Manifest option to generate RDS DDL
- D. naming scheme support with automatic partition discovery from Amazon S3

Answer: B

Explanation:

<https://docs.aws.amazon.com/emr/latest/ReleaseGuide/emr-metastore-external-hive.html>

3.What does Amazon CloudFormation provide?

- A. None of these.
- B. The ability to setup Autoscaling for Amazon EC2 instances.
- C. A template to map network resources for Amazon Web Services.
- D. A templated resource creation for Amazon Web Services.

Answer: D

4.As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon Elastic Block Store (EBS) PIOPS volume per instance and requires consistent I/O performance.

Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

- A. Ensure that the I/O block sizes for the test are randomly selected
- B. Ensure that the Amazon EBS volumes have been pre-warmed by reading all the blocks before the test
- C. Ensure that snapshots of the Amazon EBS volumes are created as a backup
- D. Ensure that the Amazon EBS volume is encrypted
- E. Ensure that the Amazon EBS volume has been pre-warmed by creating a snapshot of the volume before the test

Answer: B

5.An organization is developing a mobile social application and needs to collect logs from all devices on which it is installed. The organization is evaluating the Amazon Kinesis Data Streams to push logs and

Amazon EMR to process data. They want to store data on HDFS using the default replication factor to replicate data among the cluster, but they are concerned about the durability of the data. Currently, they are producing 300 GB of raw data daily, with additional spikes during special events. They will need to scale out the Amazon EMR cluster to match the increase in streamed data.

Which solution prevents data loss and matches compute demand?

- A. Use multiple Amazon EBS volumes on Amazon EMR to store processed data and scale out the Amazon EMR cluster as needed.
- B. Use the EMR File System and Amazon S3 to store processed data and scale out the Amazon EMR cluster as needed.
- C. Use Amazon DynamoDB to store processed data and scale out the Amazon EMR cluster as needed.
- D. use Amazon Kinesis Data Firehose and, instead of using Amazon EMR, stream logs directly into Amazon Elasticsearch Service.

Answer: B