



Free Questions for DBS-C01 by actualtestdumps

Shared by Farrell on 06-06-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

The website of a manufacturing firm makes use of an Amazon Aurora PostgreSQL database cluster.

Which settings will result in the LEAST amount of downtime for the application during failover? (Select three.)

Options:

- A- Use the provided read and write Aurora endpoints to establish a connection to the Aurora DB cluster.
- B- Create an Amazon CloudWatch alert triggering a restore in another Availability Zone when the primary Aurora DB cluster is unreachable.
- C- Edit and enable Aurora DB cluster cache management in parameter groups.
- D- Set TCP keepalive parameters to a high value.
- E- Set JDBC connection string timeout variables to a low value.
- F- Set Java DNS caching timeouts to a high value.

Answer:

A, C, E

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#AuroraPostgreSQL.BestPractices.FastF>

Question 2

Question Type: MultipleChoice

A corporation wishes to move a 1 TB Oracle database from its current location to an Amazon Aurora PostgreSQL DB cluster. The database specialist at the firm noticed that the Oracle database stores 100 GB of large binary objects (LOBs) across many tables. The Oracle database supports LOBs up to 500 MB in size and an average of 350 MB. AWS DMS was picked by the Database Specialist to transfer the data with the most replication instances.

How should the database specialist improve the transfer of the database to AWS DMS?

Options:

- A- Create a single task using full LOB mode with a LOB chunk size of 500 MB to migrate the data and LOBs together
- B- Create two tasks: task1 with LOB tables using full LOB mode with a LOB chunk size of 500 MB and task2 without LOBs
- C- Create two tasks: task1 with LOB tables using limited LOB mode with a maximum LOB size of 500 MB and task 2 without LOBs
- D- Create a single task using limited LOB mode with a maximum LOB size of 500 MB to migrate data and LOBs together

Answer:

C

Explanation:

[https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS,](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS)

'AWS DMS migrates LOB data in two phases: 1. AWS DMS creates a new row in the target table and populates the row with all data except the associated LOB value. 2.AWS DMS updates the row in the target table with the LOB data.' This means that we would need two tasks, one per phase and use limited LOB mode for best performance.

Question 3

Question Type: MultipleChoice

A database specialist has been entrusted by an ecommerce firm with designing a reporting dashboard that visualizes crucial business KPIs derived from the company's primary production database running on Amazon Aurora

a. The dashboard should be able to read data within 100 milliseconds after an update.

The Database Specialist must conduct an audit of the Aurora DB cluster's present setup and provide a cost-effective alternative. The solution must support the unexpected read demand generated by the reporting dashboard without impairing the DB cluster's write availability and performance.

Which solution satisfies these criteria?

Options:

- A-** Turn on the serverless option in the DB cluster so it can automatically scale based on demand.
- B-** Provision a clone of the existing DB cluster for the new Application team.
- C-** Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).
- D-** Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

Answer:

D

Question 4

Question Type: MultipleChoice

A Database Specialist is constructing a new Amazon Neptune DB cluster and tries to load data from Amazon S3 using the Neptune bulk loader API. The Database Specialist is confronted with the following error message:

Unable to establish a connection to the s3 endpoint. The source URL is s3:/mybucket/graphdata/ and the region code is us-east-1. Kindly confirm your Configuration S3.

Which of the following activities should the Database Specialist take to resolve the issue? (Select two.)

Options:

- A-** Check that Amazon S3 has an IAM role granting read access to Neptune
- B-** Check that an Amazon S3 VPC endpoint exists
- C-** Check that a Neptune VPC endpoint exists
- D-** Check that Amazon EC2 has an IAM role granting read access to Amazon S3
- E-** Check that Neptune has an IAM role granting read access to Amazon S3

Answer:

B, E

Explanation:

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-tutorial-IAM.html>

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html>

"An IAM role for the Neptune DB instance to assume that has an IAM policy that allows access to the data files in the S3 bucket. The policy must grant Read and List permissions." "An Amazon S3 VPC endpoint. For more information, see the Creating an Amazon S3 VPC Endpoint section."

Question 5

Question Type: MultipleChoice

A business just transitioned from an on-premises Oracle database to Amazon Aurora PostgreSQL. Following the move, the organization observed that every day around 3:00 PM, the application's response time is substantially slower. The firm has determined that the problem is with the database, not the application.

Which set of procedures should the Database Specialist do to locate the erroneous PostgreSQL query most efficiently?

Options:

- A-** Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption. Watch these dashboards during the next slow period.
- B-** Launch an Amazon EC2 instance, and install and configure an open-source PostgreSQL monitoring tool that will run reports based on the output error logs.
- C-** Modify the logging database parameter to log all the queries related to locking in the database and then check the logs after the next slow period for this information.
- D-** Enable Amazon RDS Performance Insights on the PostgreSQL database. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Answer:

D

Explanation:

<https://aws.amazon.com/blogs/database/optimizing-and-tuning-queries-in-amazon-rds-postgresql-based-on-native-and-external-tools/>

'AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load.' 'AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load.'

Question 6

Question Type: MultipleChoice

A single MySQL database was moved to Amazon Aurora by a business. The production data is stored in a database cluster in VPC PROD, whereas 12 testing environments are hosted in VPC TEST with the same AWS account. Testing has a negligible effect on the test data.

a. The development team requires that each environment be updated nightly to ensure that each test database has daily production data.

Which migration strategy will be the quickest and least expensive to implement?

Options:

- A-** Run the master in Amazon Aurora MySQL. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.
- B-** Run the master in Amazon Aurora MySQL. Take a nightly snapshot, and restore it into 12 databases in VPC_TEST using Aurora Serverless.
- C-** Run the master in Amazon Aurora MySQL. Create 12 Aurora Replicas in VPC_TEST, and script the replicas to be deleted and re-created nightly.
- D-** Run the master in Amazon Aurora MySQL using Aurora Serverless. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

Answer:

A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

Question 7

Question Type: MultipleChoice

A database professional is tasked with the task of migrating 25 GB of data files from an on-premises storage system to an Amazon Neptune database.

Which method of data loading is the FASTEST?

Options:

A- Upload the data to Amazon S3 and use the Loader command to load the data from Amazon S3 into the Neptune database.

B- Write a utility to read the data from the on-premises storage and run INSERT statements in a loop to load the data into the Neptune

database.

C- Use the AWS CLI to load the data directly from the on-premises storage into the Neptune database.

D- Use AWS DataSync to load the data directly from the on-premises storage into the Neptune database.

Answer:

A

Explanation:

1. Copy the data files to an Amazon Simple Storage Service (Amazon S3) bucket.
2. Create an IAM role with Read and List access to the bucket.
3. Create an Amazon S3 VPC endpoint.
4. Start the Neptune loader by sending a request via HTTP to the Neptune DB instance.
5. The Neptune DB instance assumes the IAM role to load the data from the bucket.

To Get Premium Files for DBS-C01 Visit

<https://www.p2pexams.com/products/dbs-c01>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/dbs-c01>

