

DUMPSBOSS.

AWS Certified Database - Specialty

Amazon AWS DBS-C01

Version Demo

Total Demo Questions: 15

Total Premium Questions: 344

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load data from Amazon S3 into the Neptune DB cluster using the Neptune bulk loader API. The Database Specialist receives the following error:

“Unable to connect to s3 endpoint. Provided source = s3://mybucket/graphdata/ and region = us-east-1. Please verify your S3 configuration.”

Which combination of actions should the Database Specialist take to troubleshoot the problem? (Choose two.)

- 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 D

Explanation:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/s3-could-not-connect-endpoint-url/>

QUESTION NO: 2

A company plans to use AWS Database Migration Service (AWS DMS) to migrate its database from one Amazon EC2 instance to another EC2 instance as a full load task. The company wants the database to be inactive during the migration. The company will use a dms.t3.medium instance to perform the migration and will use the default settings for the migration.

Which solution will MOST improve the performance of the data migration?

- A. Increase the number of tables that are loaded in parallel.
- B. Drop all indexes on the source tables.
- C. Change the processing mode from the batch optimized apply option to transactional mode.
- D. Enable Multi-AZ on the target database while the full load task is in progress.

ANSWER: B

QUESTION NO: 3

A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike. Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

- A. Enable Amazon CloudWatch Events and view the incoming T-SQL statements causing the CPU to spike.
- B. Enable Enhanced Monitoring metrics to view CPU utilization at the RDS SQL Server DB instance level.
- C. Implement a caching layer to help with repeated queries on the RDS SQL Server DB instance.
- D. Use Amazon QuickSight to view the SQL statement being run.
- E. Enable Amazon RDS Performance Insights to view the database load and filter the load by waits, SQL statements, hosts, or users.

ANSWER: B E

QUESTION NO: 4

A database professional maintains a fleet of Amazon RDS database instances that are configured to utilize the default database parameter group. A database expert must connect a custom parameter group with certain database instances.

When will the instances be allocated to this new parameter group once the database specialist performs this change?

- A. Instantaneously after the change is made to the parameter group
- B. In the next scheduled maintenance window of the DB instances
- C. After the DB instances are manually rebooted
- D. Within 24 hours after the change is made to the parameter group

ANSWER: C

QUESTION NO: 5

A financial company is hosting its web application on AWS. The application's database is hosted on Amazon RDS for MySQL with automated backups enabled. The application has caused a logical corruption of the database, which is causing the application to become unresponsive. The specific time of the corruption has been identified, and it was within the backup retention period.

How should a database specialist recover the database to the most recent point before corruption?

- A. Use the point-in-time restore capability to restore the DB instance to the specified time. No changes to the application connection string are required.
- B. Use the point-in-time restore capability to restore the DB instance to the specified time. Change the application connection string to the new, restored DB instance.

- C. Restore using the latest automated backup. Change the application connection string to the new, restored DB instance.
- D. Restore using the appropriate automated backup. No changes to the application connection string are required.

ANSWER: A

Explanation:

Reference: <https://aws.amazon.com/rds/faqs/>

QUESTION NO: 6

A company has an application that uses an Amazon DynamoDB table as its data store. During normal business days, the throughput requirements from the application are uniform and consist of 5 standard write calls per second to the DynamoDB table. Each write call has 2 KB of data.

For 1 hour each day, the company runs an additional automated job on the DynamoDB table that makes 20 write requests per second. No other application writes to the DynamoDB table. The DynamoDB table does not have to meet any additional capacity requirements.

How should a database specialist configure the DynamoDB table's capacity to meet these requirements MOST cost-effectively?

- A. Use DynamoDB provisioned capacity with 5 WCUs and auto scaling.
- B. Use DynamoDB provisioned capacity with 5 WCUs and a write-through cache that DynamoDB Accelerator (DAX) provides.
- C. Use DynamoDB provisioned capacity with 10 WCUs and auto scaling.
- D. Use DynamoDB provisioned capacity with 10 WCUs and no auto scaling.

ANSWER: A

Explanation:

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

With Application Auto Scaling, you create a *scaling policy* for a table or a global secondary index. The scaling policy specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table or index.

The scaling policy also contains a *target utilization*—the percentage of consumed provisioned throughput at a point in time. Application Auto Scaling uses a *target tracking* algorithm to adjust the provisioned throughput of the table (or index) upward or downward in response to actual workloads, so that the actual capacity utilization remains at or near your target utilization.

QUESTION NO: 7

A database specialist is designing an application to answer one-time queries. The application will query complex customer data and provide reports to end users. These reports can include many fields. The database specialist wants to give users the ability to query the database by using any of the provided fields.

The database's traffic volume will be high but variable during peak times. However, the database will not have much traffic at other times during the day.

Which solution will meet these requirements MOST cost-effectively?

- A. Amazon DynamoDB with provisioned capacity mode and auto scaling
- B. Amazon DynamoDB with on-demand capacity mode
- C. Amazon Aurora with auto scaling enabled
- D. Amazon Aurora in a serverless mode

ANSWER: B

Explanation:

Reference: <https://aws.amazon.com/about-aws/whats-new/2020/03/amazon-dynamodb-on-demand-capacity-mode-is-now-available-in-the-asia-pacific-osaka-local-region/#:~:text=On%2Ddemand%20is%20a%20flexible,to%20balance%20cost%20and%20performance>

QUESTION NO: 8

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.)

- 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

QUESTION NO: 9

A company is loading sensitive data into an Amazon Aurora MySQL database. To meet compliance requirements, the company needs to enable audit logging on the Aurora MySQL DB cluster to audit database activity. This logging will include events such as connections, disconnections, queries, and tables queried. The company also needs to publish the DB logs to Amazon CloudWatch to perform real-time data analysis.

Which solution meets these requirements?

- A.** Modify the default option group parameters to enable Advanced Auditing. Restart the database for the changes to take effect.
- B.** Create a custom DB cluster parameter group. Modify the parameters for Advanced Auditing. Modify the cluster to associate the new custom DB parameter group with the Aurora MySQL DB cluster.
- C.** Take a snapshot of the database. Create a new DB instance, and enable custom auditing and logging to CloudWatch. Deactivate the DB instance that has no logging.
- D.** Enable AWS CloudTrail for the DB instance. Create a filter that provides only connections, disconnections, queries, and tables queried.

ANSWER: B

QUESTION NO: 10

A company uses a single-node Amazon RDS for MySQL DB instance for its production database. The DB instance runs in an AWS Region in the United States.

A week before a big sales event, a new maintenance update is available for the DB instance. The maintenance update is marked as required. The company wants to minimize downtime for the DB instance and asks a database specialist to make the DB instance highly available until the sales event ends.

Which solution will meet these requirements?

- A.** Defer the maintenance update until the sales event is over.
- B.** Create a read replica with the latest update. Initiate a failover before the sales event.
- C.** Create a read replica with the latest update. Transfer all read-only traffic to the read replica during the sales event.
- D.** Convert the DB instance into a Multi-AZ deployment. Apply the maintenance update.

ANSWER: D

Explanation:

Reference: <https://aws.amazon.com/rds/features/multi-az/>

QUESTION NO: 11

A gaming company is evaluating Amazon ElastiCache as a solution to manage player leaderboards. Millions of players around the world will complete in annual tournaments. The company wants to implement an architecture that is highly

available. The company also wants to ensure that maintenance activities have minimal impact on the availability of the gaming platform.

Which combination of steps should the company take to meet these requirements? (Choose two.)

- A. Deploy an ElastiCache for Redis cluster with read replicas and Multi-AZ enabled.
- B. Deploy an ElastiCache for Memcached global datastore.
- C. Deploy a single-node ElastiCache for Redis cluster with automatic backups enabled. In the event of a failure, create a new cluster and restore data from the most recent backup.
- D. Use the default maintenance window to apply any required system changes and mandatory updates as soon as they are available.
- E. Choose a preferred maintenance window at the time of lowest usage to apply any required changes and mandatory updates.

ANSWER: A E

QUESTION NO: 12

AWS CloudFormation stack including an Amazon RDS database instance was mistakenly removed, resulting in the loss of recent data. A Database Specialist must apply RDS parameters to the CloudFormation template in order to minimize the possibility of future inadvertent instance data loss.

Which settings will satisfy this criterion? (Select three.)

- A. Set DeletionProtection to True
- B. Set MultiAZ to True
- C. Set TerminationProtection to True
- D. Set DeleteAutomatedBackups to False
- E. Set DeletionPolicy to Delete
- F. Set DeletionPolicy to Retain

ANSWER: A D F

QUESTION NO: 13

A company is using Amazon RDS for MySQL to redesign its business application. A Database Specialist has noticed that the Development team is restoring their MySQL database multiple times a day when Developers make mistakes in their schema updates. The Developers sometimes need to wait hours for the restores to complete.

Multiple team members are working on the project, making it difficult to find the correct restore point for each mistake. Which approach should the Database Specialist take to reduce downtime?

- A. Deploy multiple read replicas and have the team members make changes to separate replica instances
- B. Migrate to Amazon RDS for SQL Server, take a snapshot, and restore from the snapshot
- C. Migrate to Amazon Aurora MySQL and enable the Aurora Backtrack feature
- D. Enable the Amazon RDS for MySQL Backtrack feature

ANSWER: A

QUESTION NO: 14

A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue.

Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

- A. Grant least privilege to groups, users, and roles
- B. Allow all users to restore a database from a backup that will reduce the overall downtime to restore the database
- C. Enable multi-factor authentication for sensitive operations to access sensitive resources and API operations
- D. Use policy conditions to restrict access to selective IP addresses
- E. Use AccessList Controls policy type to restrict users for database instance deletion
- F. Enable AWS CloudTrail logging and Enhanced Monitoring

ANSWER: A C D

QUESTION NO: 15

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.)

- 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