



DBS-C01^{Q&As}

AWS Certified Database - Specialty (DBS-C01)





Pass Amazon DBS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/aws-certified-database-specialty.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

A pharmaceutical company's drug search API is using an Amazon Neptune DB cluster. A bulk uploader process automatically updates the information in the database a few times each week. A few weeks ago during a bulk upload, a database specialist noticed that the database started to respond frequently with a `ThrottlingException` error. The problem also occurred with subsequent uploads.

The database specialist must create a solution to prevent `ThrottlingException` errors for the database. The solution must minimize the downtime of the cluster.

Which solution meets these requirements?

- A. Create a read replica that uses a larger instance size than the primary DB instance. Fail over the primary DB instance to the read replica.
- B. Add a read replica to each Availability Zone. Use an instance for the read replica that is the same size as the primary DB instance. Keep the traffic between the API and the database within the Availability Zone.
- C. Create a read replica that uses a larger instance size than the primary DB instance. Offload the reads from the primary DB instance.
- D. Take the latest backup, and restore it in a DB cluster of a larger size. Point the application to the newly created DB cluster.

Correct Answer: C

Explanation: <https://docs.aws.amazon.com/neptune/latest/userguide/manage-console-add-replicas.html>

Neptune replicas connect to the same storage volume as the primary DB instance and support only read operations. Neptune replicas can offload read workloads from the primary DB instance.

QUESTION 2

A startup company in the travel industry wants to create an application that includes a personal travel assistant to display information for nearby airports based on user location. The application will use Amazon DynamoDB and must be able to access and display attributes such as airline names, arrival times, and flight numbers. However, the application must not be able to access or display pilot names or passenger counts.

Which solution will meet these requirements MOST cost-effectively?

- A. Use a proxy tier between the application and DynamoDB to regulate access to specific tables, items, and attributes.
- B. Use IAM policies with a combination of IAM conditions and actions to implement fine-grained access control.
- C. Use DynamoDB resource policies to regulate access to specific tables, items, and attributes.
- D. Configure an AWS Lambda function to extract only allowed attributes from tables based on user profiles.

Correct Answer: B

Explanation: <https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/>



QUESTION 3

A major organization maintains a number of Amazon DB clusters. Each of these clusters is configured differently to meet certain needs. These configurations may be classified into wider groups based on the team and use case.

A database administrator wishes to streamline the process of storing and updating these settings. Additionally, the database administrator want to guarantee that changes to certain configuration categories are automatically implemented to all

instances as necessary.

Which AWS service or functionality will assist in automating and achieving this goal?

- A. AWS Systems Manager Parameter Store
- B. DB parameter group
- C. AWS Config
- D. AWS Secrets Manager

Correct Answer: B

Explanation: Database parameters specify how the database is configured. For example, database parameters can specify the amount of resources, such as memory, to allocate to a database.

QUESTION 4

A company runs a customer relationship management (CRM) system that is hosted on- premises with a MySQL database as the backend. A custom stored procedure is used to send email notifications to another system when data is inserted into a table. The company has noticed that the performance of the CRM system has decreased due to database reporting applications used by various teams. The company requires an AWS solution that would reduce maintenance, improve performance, and accommodate the email notification feature.

Which AWS solution meets these requirements?

- A. Use MySQL running on an Amazon EC2 instance with Auto Scaling to accommodate the reporting applications. Configure a stored procedure and an AWS Lambda function that uses Amazon SES to send email notifications to the other system.
- B. Use Amazon Aurora MySQL in a multi-master cluster to accommodate the reporting applications. Configure Amazon RDS event subscriptions to publish a message to an Amazon SNS topic and subscribe the other system's email address to the topic.
- C. Use MySQL running on an Amazon EC2 instance with a read replica to accommodate the reporting applications. Configure Amazon SES integration to send email notifications to the other system.
- D. Use Amazon Aurora MySQL with a read replica for the reporting applications. Configure a stored procedure and an AWS Lambda function to publish a message to an Amazon SNS topic. Subscribe the other system's email address to the topic.

Correct Answer: D



RDS event subscriptions do not cover "data is inserted into a table" - see https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_Events.Messages.html We can use stored procedure to invoke Lambda function - <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html>

QUESTION 5

A company migrated one of its business-critical database workloads to an Amazon Aurora Multi-AZ DB cluster. The company requires a very low RTO and needs to improve the application recovery time after database failovers.

Which approach meets these requirements?

- A. Set the max_connections parameter to 16,000 in the instance-level parameter group.
- B. Modify the client connection timeout to 300 seconds.
- C. Create an Amazon RDS Proxy database proxy and update client connections to point to the proxy endpoint.
- D. Enable the query cache at the instance level.

Correct Answer: C

Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up

to 66% and database credentials, authentication, and access can be managed through integration with AWS Secrets Manager and AWS Identity and Access Management (IAM).

<https://aws.amazon.com/rds/proxy/>

[DBS-C01 VCE Dumps](#)

[DBS-C01 Practice Test](#)

[DBS-C01 Exam Questions](#)