



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.pass4itsure.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

For the first time, a database professional is establishing a test graph database on Amazon Neptune. The database expert must input millions of rows of test observations from an Amazon S3.csv file. The database professional uploaded the data to the Neptune DB instance through a series of API calls.

Which sequence of actions enables the database professional to upload the data most quickly? (Select three.)

- A. Ensure Amazon Cognito returns the proper AWS STS tokens to authenticate the Neptune DB instance to the S3 bucket hosting the CSV file.
- B. Ensure the vertices and edges are specified in different .csv files with proper header column formatting.
- C. Use AWS DMS to move data from Amazon S3 to the Neptune Loader.
- D. Curl the S3 URI while inside the Neptune DB instance and then run the addVertex or addEdge commands.
- E. Ensure an IAM role for the Neptune DB instance is configured with the appropriate permissions to allow access to the file in the S3 bucket.
- F. Create an S3 VPC endpoint and issue an HTTP POST to the databaseTMs loader endpoint.

Correct Answer: BEF

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

QUESTION 2

A company is developing a new web application. An AWS CloudFormation template was created as a part of the build process.

Recently, a change was made to an AWS::RDS::DBInstance resource in the template. The CharacterSetName property was changed to allow the application to process international text. A change set was generated using the new template, which indicated that the existing DB instance should be replaced during an upgrade.

What should a database specialist do to prevent data loss during the stack upgrade?

- A. Create a snapshot of the DB instance. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot. Update the stack.
- B. Modify the stack policy using the aws cloudformation update-stack command and the set-stack-policy command, then make the DB resource protected.
- C. Create a snapshot of the DB instance. Update the stack. Restore the database to a new instance.
- D. Deactivate any applications that are using the DB instance. Create a snapshot of the DB instance. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot. Update the stack and reactivate the applications.

Correct Answer: D

To preserve your data, perform the following procedure:



1. Deactivate any applications that are using the DB instance so that there's no activity on the DB instance.
2. Create a snapshot of the DB instance. For more information about creating DB snapshots
3. If you want to restore your instance using a DB snapshot, modify the updated template with your DB instance changes and add the DBSnapshotIdentifier property with the ID of the DB snapshot that you want to use
4. Update the stack.

Reference: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-rds-database-instance.html>

QUESTION 3

A company is using 5 TB Amazon RDS DB instances and needs to maintain 5 years of monthly database backups for compliance purposes. A Database Administrator must provide Auditors with data within 24 hours. Which solution will meet these requirements and is the MOST operationally efficient?

- A. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot. Move the snapshot to the company's Amazon S3 bucket.
- B. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot.
- C. Create an RDS snapshot schedule from the AWS Management Console to take a snapshot every 30 days.
- D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Correct Answer: A

Explanation: Unlike automated backups, manual snapshots aren't subject to the backup retention period. Snapshots don't expire. For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html

QUESTION 4

A software-as-a-service (SaaS) company is using an Amazon Aurora Serverless DB cluster for its production MySQL database. The DB cluster has general logs and slow query logs enabled. A database engineer must use the most operationally efficient solution with minimal resource utilization to retain the logs and facilitate interactive search and analysis.

Which solution meets these requirements?

- A. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- B. Download the logs from the DB cluster and store them in Amazon S3 by using manual scripts. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- C. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket. Use Amazon Elasticsearch Service (Amazon ES) and Kibana to search and analyze the logs.
- D. Use Amazon CloudWatch Logs Insights to search and analyze the logs when the logs are automatically uploaded by



the DB cluster.

Correct Answer: D

Explanation: <https://aws.amazon.com/premiumsupport/knowledge-center/aurora-serverless-logs-enable-view/>
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.html>

QUESTION 5

A company is developing a multi-tier web application hosted on AWS using Amazon Aurora as the database. The application needs to be deployed to production and other non-production environments. A Database Specialist needs to specify different MasterUsername and MasterUserPassword properties in the AWS CloudFormation templates used for automated deployment. The CloudFormation templates are version controlled in the company's code repository. The company also needs to meet compliance requirement by routinely rotating its database master password for production.

What is most secure solution to store the master password?

- A. Store the master password in a parameter file in each environment. Reference the environment-specific parameter file in the CloudFormation template.
- B. Encrypt the master password using an AWS KMS key. Store the encrypted master password in the CloudFormation template.
- C. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.
- D. Use the ssm dynamic reference to retrieve the master password stored in the AWS Systems Manager Parameter Store and enable automatic rotation.

Correct Answer: C

"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code. Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines." <https://aws.amazon.com/blogs/mt/using-aws-systems-manager-parameter-store-secure-string-parameters-in-aws-cloudformation-templates/>
<https://aws.amazon.com/blogs/security/how-to-use-aws-secretsmanager-rotate-credentials-amazon-rds-database-types-oracle/>

[Latest DBS-C01 Dumps](#)

[DBS-C01 Exam Questions](#)

[DBS-C01 Braindumps](#)