



Vendor: Amazon

Exam Code: DBS-C01

Exam Name: AWS Certified Database - Specialty (DBS-C01) Exam

Version: DEMO

QUESTION 1

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

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.Performance

For a full load task, we recommend that you drop primary key indexes, secondary indexes, referential integrity constraints, and data manipulation language (DML) triggers. Or you can delay their creation until after the full load tasks are complete. You don't need indexes during a full load task, and indexes incur maintenance overhead if they are present. Because the full load task loads groups of tables at a time, referential integrity constraints are violated. Similarly, insert, update, and delete triggers can cause errors, for example if a row insert is triggered for a previously bulk loaded table. Other types of triggers also affect performance due to added processing.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html

QUESTION 2

A company is running a business-critical application on premises by using Microsoft SQL Server. A database specialist is planning to migrate the instance with several databases to the AWS Cloud. The database specialist will use SQL Server Standard edition hosted on Amazon EC2 Windows instances. The solution must provide high availability and must avoid a single point of failure in the SQL Server deployment architecture. Which solution will meet these requirements?

- A. Create Amazon RDS for SQL Server Multi-AZ DB instances. Use Amazon S3 as a shared storage option to host the databases.
- B. Set up Always On Failover Cluster Instances as a single SQL Server instance. Use Multi-AZ Amazon FSx for Windows File Server as a shared storage option to host the databases.
- C. Set up Always On availability groups to group one or more user databases that fail over together across multiple SQL Server instances. Use Multi-AZ Amazon FSx for Windows File Server as a shared storage option to host the databases.
- D. Create an Application Load Balancer to distribute database traffic across multiple EC2 instances in multiple Availability Zones. Use Amazon S3 as a shared storage option to host the databases.

Answer: B

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/migration-sql-server/ec2-fci.html>

An FCI is generally preferable over an Always on availability group when: You're using SQL Server Standard edition instead of Enterprise edition.

QUESTION 3

A company is launching a new Amazon RDS for MySQL Multi-AZ DB instance to be used as a data store for a custom-built application. After a series of tests with point-in-time recovery

disabled, the company decides that it must have point-in-time recovery reenabled before using the DB instance to store production data.

What should a database specialist do so that point-in-time recovery can be successful?

- A. Enable binary logging in the DB parameter group used by the DB instance.
- B. Modify the DB instance and enable audit logs to be pushed to Amazon CloudWatch Logs.
- C. Modify the DB instance and configure a backup retention period
- D. Set up a scheduled job to create manual DB instance snapshots.

Answer: C

Explanation:

You can restore a DB instance to a specific point in time (PITR), creating a new DB instance. To support PITR, your DB instances must have backup retention set to a nonzero value.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/custom-backup-sqlserver.html>

<https://aws.amazon.com/blogs/database/setting-up-a-binlog-server-for-amazon-rds-mysql-and-mariadb-using-mariadb-maxscale/>

"After you run the command, it's okay to enable backup retention on the RDS instance by using the AWS CLI or the console. Enabling backup retention also enables binary logging."

<https://aws.amazon.com/blogs/storage/point-in-time-recovery-and-continuous-backup-for-amazon-rds-with-aws-backup/>

QUESTION 4

A database specialist wants to ensure that an Amazon Aurora DB cluster is always automatically upgraded to the most recent minor version available. Noticing that there is a new minor version available, the database specialist has issues an AWS CLI command to enable automatic minor version updates. The command runs successfully, but checking the Aurora DB cluster indicates that no update to the Aurora version has been made.

What might account for this? (Choose two.)

- A. The new minor version has not yet been designated as preferred and requires a manual upgrade.
- B. Configuring automatic upgrades using the AWS CLI is not supported. This must be enabled expressly using the AWS Management Console.
- C. Applying minor version upgrades requires sufficient free space.
- D. The AWS CLI command did not include an apply-immediately parameter.
- E. Aurora has detected a breaking change in the new minor version and has automatically rejected the upgrade.

Answer: AD

Explanation:

When Amazon RDS designates a minor engine version as the preferred minor engine version, each database that meets both of the following conditions is upgraded to the minor engine version automatically.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Upgrading.html

Call the modify-db-instance Amazon CLI command. Specify the name of your DB instance for the --db-instance-identifier option and true for the --auto-minor-version-upgrade option. Optionally, specify the --apply-immediately option to immediately enable this setting for your DB instance. Run a separate modify-db-instance command for each DB instance in the cluster.

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Updates.Patching.html#AuroraMySQL.Updates.AMVU

QUESTION 5

A company wants to improve its ecommerce website on AWS. A database specialist decides to add Amazon ElastiCache for Redis in the implementation stack to ease the workload off the database and shorten the website response times. The database specialist must also ensure the ecommerce website is highly available within the company's AWS Region.

How should the database specialist deploy ElastiCache to meet this requirement?

- A. Launch an ElastiCache for Redis cluster using the AWS CLI with the -cluster-enabled switch.
- B. Launch an ElastiCache for Redis cluster and select read replicas in different Availability Zones.
- C. Launch two ElastiCache for Redis clusters in two different Availability Zones. Configure Redis streams to replicate the cache from the primary cluster to another.
- D. Launch an ElastiCache cluster in the primary Availability Zone and restore the cluster's snapshot to a different Availability Zone during disaster recovery.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html>

You can enable Multi-AZ only on Redis (cluster mode disabled) clusters that have at least one available read replica. Clusters without read replicas do not provide high availability or fault tolerance.

QUESTION 6

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.

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 7

A financial company is running an Amazon Redshift cluster for one of its data warehouse solutions. The company needs to generate connection logs, user logs, and user activity logs. The company also must make these logs available for future analysis.

Which combination of steps should a database specialist take to meet these requirements?

(Choose two.)

- A. Edit the database configuration of the cluster by enabling audit logging. Direct the logging to a specified log group in Amazon CloudWatch Logs.
- B. Edit the database configuration of the cluster by enabling audit logging. Direct the logging to a specified Amazon S3 bucket
- C. Modify the cluster by enabling continuous delivery of AWS CloudTrail logs to Amazon S3.
- D. Create a new parameter group with the enable_user_activity_logging parameter set to true. Configure the cluster to use the new parameter group.
- E. Modify the system table to enable logging for each user.

Answer: AD

Explanation:

AWS CloudWatch Logs are stored indefinitely and CloudWatch Log Insights is used to analyze the logs and query upon them.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.html>

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

"Log retention - By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period between 10 years and one day."

<https://docs.aws.amazon.com/redshift/latest/mgmt/db-auditing.html>

QUESTION 8

A company's database specialist is building an Amazon RDS for Microsoft SQL Server DB instance to store hundreds of records in CSV format. A customer service tool uploads the records to an Amazon S3 bucket.

An employee who previously worked at the company already created a custom stored procedure to map the necessary CSV fields to the database tables. The database specialist needs to implement a solution that reuses this previous work and minimizes operational overhead.

Which solution will meet these requirements?

- A. Create an Amazon S3 event to invoke an AWS Lambda function. Configure the Lambda function to parse the .csv file and use a SQL client library to run INSERT statements to load the data into the tables.
- B. Write a custom .NET app that is hosted on Amazon EC2. Configure the .NET app to load the .csv file and call the custom stored procedure to insert the data into the tables.
- C. Download the .csv file from Amazon S3 to the RDS D drive by using an AWS msdb stored procedure. Call the custom stored procedure to insert the data from the RDS D drive into the tables.
- D. Create an Amazon S3 event to invoke AWS Step Functions to parse the .csv file and call the custom stored procedure to insert the data into the tables.

Answer: C

Explanation:

<https://www.mssqltips.com/sqlservertip/6619/rds-sql-server-data-import-from-amazon-s3/>

Amazon Web Service (AWS) recently announced a new feature of its Relational Database Service (RDS) for SQL Server. This feature allows a native integration between Amazon RDS SQL Server and Amazon S3. With this integration, it's now possible to import files from an Amazon S3 bucket into a local folder of the RDS instance. Similarly, files from that folder can be exported to S3. The RDS local folder path is D:\S3\.

QUESTION 9

A company is running a blogging platform. A security audit determines that the Amazon RDS DB instance that is used by the platform is not configured to encrypt the data at rest. The company

must encrypt the DB instance within 30 days.

What should a database specialist do to meet this requirement with the LEAST amount of downtime?

- A. Create a read replica of the DB instance, and enable encryption. When the read replica is available, promote the read replica and update the endpoint that is used by the application. Delete the unencrypted DB instance.
- B. Take a snapshot of the DB instance. Make an encrypted copy of the snapshot. Restore the encrypted snapshot. When the new DB instance is available, update the endpoint that is used by the application. Delete the unencrypted DB instance.
- C. Create a new encrypted DB instance. Perform an initial data load, and set up logical replication between the two DB instances. When the new DB instance is in sync with the source DB instance, update the endpoint that is used by the application. Delete the unencrypted DB instance.
- D. Convert the DB instance to an Amazon Aurora DB cluster, and enable encryption. When the DB cluster is available, update the endpoint that is used by the application to the cluster endpoint. Delete the unencrypted DB instance.

Answer: C

Explanation:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresql-db-instance.html>

When the new, encrypted copy of the DB instance becomes available, you can point your applications to the new database. However, if your project doesn't allow for significant downtime for this activity, you need an alternate approach that helps minimize the downtime. This pattern uses the AWS Database Migration Service (AWS DMS) to migrate and continuously replicate the data so that the cutover to the new, encrypted database can be done with minimal downtime.

QUESTION 10

A database specialist wants to ensure that an Amazon Aurora DB cluster is always automatically upgraded to the most recent minor version available. Noticing that there is a new minor version available, the database specialist has issues an AWS CLI command to enable automatic minor version updates. The command runs successfully, but checking the Aurora DB cluster indicates that no update to the Aurora version has been made.

What might account for this? (Choose two.)

- A. The new minor version has not yet been designated as preferred and requires a manual upgrade.
- B. Configuring automatic upgrades using the AWS CLI is not supported. This must be enabled expressly using the AWS Management Console.
- C. Applying minor version upgrades requires sufficient free space.
- D. The AWS CLI command did not include an apply-immediately parameter.
- E. Aurora has detected a breaking change in the new minor version and has automatically rejected the upgrade.

Answer: AD

Explanation:

When Amazon RDS designates a minor engine version as the preferred minor engine version, each database that meets both of the following conditions is upgraded to the minor engine version automatically.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Upgrading.html

QUESTION 11

A software company uses an Amazon RDS for MySQL Multi-AZ DB instance as a data store for its critical applications. During an application upgrade process, a database specialist runs a custom SQL script that accidentally removes some of the default permissions of the master user.

What is the MOST operationally efficient way to restore the default permissions of the master user?

- A. Modify the DB instance and set a new master user password.
- B. Use AWS Secrets Manager to modify the master user password and restart the DB instance.
- C. Create a new master user for the DB instance.
- D. Review the IAM user that owns the DB instance, and add missing permissions.

Answer: A

Explanation:

If you accidentally delete the permissions for the master user, you can restore them by modifying the DB instance and setting a new master user password.

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.MasterAccounts.html>

QUESTION 12

A company is launching a new Amazon RDS for MySQL Multi-AZ DB instance to be used as a data store for a custom-built application. After a series of tests with point-in-time recovery disabled, the company decides that it must have point-in-time recovery reenabled before using the DB instance to store production data.

What should a database specialist do so that point-in-time recovery can be successful?

- A. Enable binary logging in the DB parameter group used by the DB instance.
- B. Modify the DB instance and enable audit logs to be pushed to Amazon CloudWatch Logs.
- C. Modify the DB instance and configure a backup retention period
- D. Set up a scheduled job to create manual DB instance snapshots.

Answer: C

Explanation:

You can restore a DB instance to a specific point in time (PITR), creating a new DB instance. To support PITR, your DB instances must have backup retention set to a nonzero value.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/custom-backup-sqlserver.html>

QUESTION 13

A company has a 4 TB on-premises Oracle Real Application Clusters (RAC) database. The company wants to migrate the database to AWS and reduce licensing costs. The company's application team wants to store JSON payloads that expire after 28 hours. The company has development capacity if code changes are required.

Which solution meets these requirements?

- A. Use Amazon DynamoDB and leverage the Time to Live (TTL) feature to automatically expire the data.
- B. Use Amazon RDS for Oracle with Multi-AZ. Create an AWS Lambda function to purge the expired data. Schedule the Lambda function to run daily using Amazon EventBridge.

- C. Use Amazon DocumentDB with a read replica in a different Availability Zone. Use DocumentDB change streams to expire the data.
- D. Use Amazon Aurora PostgreSQL with Multi-AZ and leverage the Time to Live (TTL) feature to automatically expire the data.

Answer: A

Explanation:

Amazon DynamoDB Time to Live (TTL) allows you to define a per-item timestamp to determine when an item is no longer needed. Shortly after the date and time of the specified timestamp, DynamoDB deletes the item from your table without consuming any write throughput.

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

Amazon DynamoDB Time to Live (TTL) allows you to define a per-item timestamp to determine when an item is no longer needed. Shortly after the date and time of the specified timestamp, DynamoDB deletes the item from your table without consuming any write throughput. TTL is provided at no extra cost as a means to reduce stored data volumes by retaining only the items that remain current for your workload's needs.

TTL is useful if you store items that lose relevance after a specific time. The following are example TTL use cases:

- Remove user or sensor data after one year of inactivity in an application.
- Archive expired items to an Amazon S3 data lake via Amazon DynamoDB Streams and AWS Lambda.
- Retain sensitive data for a certain amount of time according to contractual or regulatory obligations.

QUESTION 14

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

Explanation:

For a full load task, we recommend that you drop primary key indexes, secondary indexes, referential integrity constraints, and data manipulation language (DML) triggers. Or you can delay their creation until after the full load tasks are complete. You don't need indexes during a full load task, and indexes incur maintenance overhead if they are present. Because the full load task loads groups of tables at a time, referential integrity constraints are violated. Similarly, insert, update, and delete triggers can cause errors, for example if a row insert is triggered for a previously bulk loaded table. Other types of triggers also affect performance due to added processing.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html

QUESTION 15

A vehicle insurance company needs to choose a highly available database to track vehicle owners and their insurance details. The persisted data should be immutable in the database, including the complete and sequenced history of changes over time with all the owners and insurance transfer details for a vehicle. The data should be easily verifiable for the data lineage of an insurance claim.

Which approach meets these requirements with MINIMAL effort?

- A. Create a blockchain to store the insurance details. Validate the data using a hash function to verify the data lineage of an insurance claim.
- B. Create an Amazon DynamoDB table to store the insurance details. Validate the data using AWS DMS validation by moving the data to Amazon S3 to verify the data lineage of an insurance claim.
- C. Create an Amazon QLDB ledger to store the insurance details. Validate the data by choosing the ledger name in the digest request to verify the data lineage of an insurance claim.
- D. Create an Amazon Aurora database to store the insurance details. Validate the data using AWS DMS validation by moving the data to Amazon S3 to verify the data lineage of an insurance claim.

Answer: C

Explanation:

QLDB API

You can also verify a document revision by using the Amazon QLDB API with an AWS SDK or the AWS CLI. The QLDB API provides the following operations for use by application programs:

- **GetDigest** – Returns the digest of a ledger at the latest committed block in the journal. The response includes a 256-bit hash value and a block address.
- **GetBlock** – Returns a block object at a specified address in a journal. Also returns a proof of the specified block for verification if **DigestTipAddress** is provided.
- **GetRevision** – Returns a revision data object for a specified document ID and block address. Also returns a proof of the specified revision for verification if **DigestTipAddress** is provided.

<https://docs.aws.amazon.com/qldb/latest/developerguide/verification.verify.html>

QUESTION 16

A company has deployed an application that uses an Amazon RDS for MySQL DB cluster. The DB cluster uses three read replicas. The primary DB instance is an 8XL-sized instance, and the read replicas are each XL-sized instances.

Users report that database queries are returning stale data. The replication lag indicates that the replicas are 5 minutes behind the primary DB instance. Status queries on the replicas show that the SQL_THREAD is 10 binlogs behind the IO_THREAD and that the IO_THREAD is 1 binlog behind the primary.

Which changes will reduce the lag? (Choose two.)

- A. Deploy two additional read replicas matching the existing replica DB instance size.
- B. Migrate the primary DB instance to an Amazon Aurora MySQL DB cluster and add three Aurora Replicas.
- C. Move the read replicas to the same Availability Zone as the primary DB instance.
- D. Increase the instance size of the primary DB instance within the same instance class.
- E. Increase the instance size of the read replicas to the same size and class as the primary DB instance.

Answer: BE

Explanation:

Read replicas should be of the same size and class as primary DB.

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-mysql-high-replica-lag/>

Aurora seems to be better in handling the lags.

<https://www.quora.com/What-is-the-difference-between-a-RDS-read-replica-and-an-Aurora-Aurora-Read-replica>

QUESTION 17

A company's applications store data in Amazon Aurora MySQL DB clusters. The company has separate AWS accounts for its production, test, and development environments. To test new functionality in the test environment, the company's development team requires a copy of the production database four times a day.

Which solution meets this requirement with the MOST operational efficiency?

- A. Take a manual snapshot in the production account. Share the snapshot with the test account. Restore the database from the snapshot.
- B. Take a manual snapshot in the production account. Export the snapshot to Amazon S3. Copy the snapshot to an S3 bucket in the test account. Restore the database from the snapshot.
- C. Share the Aurora DB cluster with the test account. Create a snapshot of the production database in the test account. Restore the database from the snapshot.
- D. Share the Aurora DB cluster with the test account. Create a clone of the production database in the test account.

Answer: D

Explanation:

Creating and restoring a database snapshot. You can create a clone of one of your Aurora DB clusters and share the clone.

Reference:

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

QUESTION 18

A company is migrating a database in an Amazon RDS for SQL Server DB instance from one AWS Region to another. The company wants to minimize database downtime during the migration.

Which strategy should the company choose for this cross-Region migration?

- A. Back up the source database using native backup to an Amazon S3 bucket in the same Region. Then restore the backup in the target Region.
- B. Back up the source database using native backup to an Amazon S3 bucket in the same Region. Use Amazon S3 Cross-Region Replication to copy the backup to an S3 bucket in the target Region. Then restore the backup in the target Region.
- C. Configure AWS Database Migration Service (AWS DMS) to replicate data between the source and the target databases. Once the replication is in sync, terminate the DMS task.
- D. Add an RDS for SQL Server cross-Region read replica in the target Region. Once the replication is in sync, promote the read replica to master.

Answer: C

Explanation:

Amazon RDS supports native backup and restore for Microsoft SQL Server databases using full backup files (.bak files). When you use RDS, you access files stored in Amazon S3 rather than using the local file system on the database server.

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html>

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



10% Discount Coupon Code: ASTR14