

➤ **Vendor: Amazon**

➤ **Exam Code: DBS-C01**

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

➤ **New Updated Questions from [Braindump2go](#)**

➤ **(Updated in [February/2022](#))**

## [Visit Braindump2go and Download Full Version DBS-C01 Exam Dumps](#)

### QUESTION 168

A company is setting up a new Amazon RDS for SQL Server DB instance. The company wants to enable SQL Server auditing on the database.

Which combination of steps should a database specialist take to meet this requirement? (Choose two.)

- A. Create a service-linked role for Amazon RDS that grants permissions for Amazon RDS to store audit logs on Amazon S3.
- B. Set up a parameter group to configure an IAM role and an Amazon S3 bucket for audit log storage. Associate the parameter group with the DB instance.
- C. Disable Multi-AZ on the DB instance, and then enable auditing. Enable Multi-AZ after auditing is enabled.
- D. Disable automated backup on the DB instance, and then enable auditing. Enable automated backup after auditing is enabled.
- E. Set up an options group to configure an IAM role and an Amazon S3 bucket for audit log storage. Associate the options group with the DB instance.

**Answer: AE**

#### **Explanation:**

To do this, you create an IAM role and delegate permissions so that the Amazon RDS service can use your Amazon S3 bucket.

RDS uploads the completed audit logs to your S3 bucket, using the IAM role that you provide. If you enable retention, RDS keeps your audit logs on your DB instance for the configured period of time.

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.Options.Audit.html>

### QUESTION 169

A database specialist is creating an AWS CloudFormation stack. The database specialist wants to prevent accidental deletion of an Amazon RDS ProductionDatabase resource in the stack.

Which solution will meet this requirement?

- A. Create a stack policy to prevent updates. Include "Effect" : "ProductionDatabase" and "Resource" : "Deny" in the policy.
- B. Create an AWS CloudFormation stack in XML format. Set xAttribute as false.
- C. Create an RDS DB instance without the DeletionPolicy attribute. Disable termination protection.
- D. Create a stack policy to prevent updates. Include "Effect" : "Deny" and "Resource" : "ProductionDatabase" in the policy.

**Answer: D**

#### **Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-accidental-updates/>

[DBS-C01 Exam Dumps](#) [DBS-C01 Exam Questions](#) [DBS-C01 PDF Dumps](#) [DBS-C01 VCE Dumps](#)

<https://www.braindump2go.com/dbs-c01.html>

**QUESTION 170**

An ecommerce company migrates an on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility). After the migration, a database specialist realizes that encryption at rest has not been turned on for the Amazon DocumentDB cluster.

What should the database specialist do to enable encryption at rest for the Amazon DocumentDB cluster?

- A. Take a snapshot of the Amazon DocumentDB cluster. Restore the unencrypted snapshot as a new cluster while specifying the encryption option, and provide an AWS Key Management Service (AWS KMS) key.
- B. Enable encryption for the Amazon DocumentDB cluster on the AWS Management Console. Reboot the cluster.
- C. Modify the Amazon DocumentDB cluster by using the modify-db-cluster command with the --storage-encrypted parameter set to true.
- D. Add a new encrypted instance to the Amazon DocumentDB cluster, and then delete an unencrypted instance from the cluster. Repeat until all instances are encrypted.

**Answer: B**

**Explanation:**

You can enable or disable encryption at rest on an Amazon DocumentDB cluster when the cluster is provisioned using either the AWS Management Console.

Reference: <https://docs.aws.amazon.com/documentdb/latest/developerguide/encryption-at-rest>

**QUESTION 171**

A company that analyzes the stock market has two offices: one in the us-east-1 Region and another in the eu-west-2 Region. The company wants to implement an AWS database solution that can provide fast and accurate updates. The office in eu-west-2 has dashboards with complex analytical queries to display the data. The company will use these dashboards to make buying decisions, so the dashboards must have access to the application data in less than 1 second.

Which solution meets these requirements and provides the MOST up-to-date dashboard?

- A. Deploy an Amazon RDS DB instance in us-east-1 with a read replica instance in eu-west-2. Create an Amazon ElastiCache cluster in eu-west-2 to cache data from the read replica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in us-east-1 with replication into eu-west-2. Use multi-active replication to ensure that updates are quickly propagated to eu-west-2.
- C. Use an Amazon Aurora global database. Deploy the primary DB cluster in us-east-1. Deploy the secondary DB cluster in eu-west-2. Configure the dashboard application to read from the secondary cluster.
- D. Deploy an Amazon RDS for MySQL DB instance in us-east-1 with a read replica instance in eu-west-2. Configure the dashboard application to read from the read replica.

**Answer: C**

**QUESTION 172**

A company is running its customer feedback application on Amazon Aurora MySQL. The company runs a report every day to extract customer feedback, and a team reads the feedback to determine if the customer comments are positive or negative. It sometimes takes days before the company can contact unhappy customers and take corrective measures. The company wants to use machine learning to automate this workflow.

Which solution meets this requirement with the LEAST amount of effort?

- A. Export the Aurora MySQL database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Use Amazon Comprehend to run sentiment analysis on the exported files.
- B. Export the Aurora MySQL database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Use Amazon Sage
- C. Set up Aurora native integration with Amazon Comprehend. Use SQL functions to extract sentiment analysis.

- D. Set up Aurora native integration with Amazon SageMaker. Use SQL functions to extract sentiment analysis.

**Answer: C**

**Explanation:**

Amazon Comprehend is a natural-language processing (NLP) service that uses machine learning to uncover valuable insights and connections in text.

Reference: <https://aws.amazon.com/comprehend>

#### **QUESTION 173**

A bank plans to use an Amazon RDS for MySQL DB instance. The database should support read-intensive traffic with very few repeated queries.

Which solution meets these requirements?

- A. Create an Amazon ElastiCache cluster. Use a write-through strategy to populate the cache.
- B. Create an Amazon ElastiCache cluster. Use a lazy loading strategy to populate the cache.
- C. Change the DB instance to Multi-AZ with a standby instance in another AWS Region.
- D. Create a read replica of the DB instance. Use the read replica to distribute the read traffic.

**Answer: D**

#### **QUESTION 174**

A database specialist has a fleet of Amazon RDS DB instances that use the default DB parameter group. The database specialist needs to associate a custom parameter group with some of the DB instances.

After the database specialist makes this change, when will the instances be assigned to this new parameter group?

- 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**

**Explanation:**

To apply the latest parameter changes to that DB instance, manually reboot the DB instance.

Reference: [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_WorkingWithParamGroups.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithParamGroups.html)

#### **QUESTION 175**

A company is planning on migrating a 500-GB database from Oracle to Amazon Aurora PostgreSQL using the AWS Schema Conversion Tool (AWS SCT) and AWS DMS. The database does not have any stored procedures to migrate but has some tables that are large or partitioned. The application is critical for business so a migration with minimal downtime is preferred.

Which combination of steps should a database specialist take to accelerate the migration process? (Choose three.)

- A. Use the AWS SCT data extraction agent to migrate the schema from Oracle to Aurora PostgreSQL.
- B. For the large tables, change the setting for the maximum number of tables to load in parallel and perform a full load using AWS DMS.
- C. For the large tables, create a table settings rule with a parallel load option in AWS DMS, then perform a full load using DMS.
- D. Use AWS DMS to set up change data capture (CDC) for continuous replication until the cutover date.
- E. Use AWS SCT to convert the schema from Oracle to Aurora PostgreSQL.
- F. Use AWS DMS to convert the schema from Oracle to Aurora PostgreSQL and for continuous replication.

**Answer: BF**

**QUESTION 176**

A company is migrating an IBM Informix database to a Multi-AZ deployment of Amazon RDS for SQL Server with Always On Availability Groups (AGs). SQL Server Agent jobs on the Always On AG listener run at 5-minute intervals to synchronize data between the Informix database and the SQL Server database. Users experience hours of stale data after a successful failover to the secondary node with minimal latency.

What should a database specialist do to ensure that users see recent data after a failover?

- A. Set TTL to less than 30 seconds for cached DNS values on the Always On AG listener.
- B. Break up large transactions into multiple smaller transactions that complete in less than 5 minutes.
- C. Set the databases on the secondary node to read-only mode.
- D. Create the SQL Server Agent jobs on the secondary node from a script when the secondary node takes over after a failure.

**Answer: C**

**Explanation:**

After a failover, client applications that need to access the primary databases must connect to the new primary replica. Also, if the new secondary replica is configured to allow read-only access, read-only client applications can connect to it.

Reference: <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/failover-and-failover-modes-always-on-availability-groups?view=sql-server-ver15>

**QUESTION 177**

A database specialist needs to configure an Amazon RDS for MySQL DB instance to close non-interactive connections that are inactive after 900 seconds.

What should the database specialist do to accomplish this task?

- A. Create a custom DB parameter group and set the wait\_timeout parameter value to 900. Associate the DB instance with the custom parameter group.
- B. Connect to the MySQL database and run the SET SESSION wait\_timeout=900 command.
- C. Edit the my.cnf file and set the wait\_timeout parameter value to 900. Restart the DB instance.
- D. Modify the default DB parameter group and set the wait\_timeout parameter value to 900.

**Answer: B**

**Explanation:**

If we set the wait\_timeout variable for a session, it will valid only for a particular session. But when we set the wait\_timeout variable globally it will valid for all the sessions.

Reference: <https://dilsichandrasena.medium.com/changing-mysql-wait-timeout-variable-f16ebed1efce>

**QUESTION 178**

A company is running its production databases in a 3 TB Amazon Aurora MySQL DB cluster. The DB cluster is deployed to the us-east-1 Region. For disaster recovery (DR) purposes, the company's database specialist needs to make the DB cluster rapidly available in another AWS Region to cover the production load with an RTO of less than 2 hours.

What is the MOST operationally efficient solution to meet these requirements?

- A. Implement an AWS Lambda function to take a snapshot of the production DB cluster every 2 hours, and copy that snapshot to an Amazon S3 bucket in the DR Region. Restore the snapshot to an appropriately sized DB cluster in the DR Region.
- B. Add a cross-Region read replica in the DR Region with the same instance type as the current primary instance. If the read replica in the DR Region needs to be used for production, promote the read replica to become a standalone DB cluster.
- C. Create a smaller DB cluster in the DR Region. Configure an AWS Database Migration Service (AWS DMS) task with change data capture (CDC) enabled to replicate data from the current production DB cluster to the DB cluster in the DR Region.
- D. Create an Aurora global database that spans two Regions. Use AWS Database Migration Service

**[DBS-C01 Exam Dumps](#) [DBS-C01 Exam Questions](#) [DBS-C01 PDF Dumps](#) [DBS-C01 VCE Dumps](#)**

**<https://www.braindump2go.com/dbs-c01.html>**

(AWS DMS) to migrate the existing database to the new global database.

**Answer: B**

**Explanation:**

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_ReadRepl.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html)

#### **QUESTION 179**

A company has an on-premises SQL Server database. The users access the database using Active Directory authentication. The company successfully migrated its database to Amazon RDS for SQL Server. However, the company is concerned about user authentication in the AWS Cloud environment. Which solution should a database specialist provide for the user to authenticate?

- A. Deploy Active Directory Federation Services (AD FS) on premises and configure it with an on-premises Active Directory. Set up delegation between the on-premises AD FS and AWS Security Token Service (AWS STS) to map user identities to a role using theAmazonRDSDirectoryServiceAccess managed IAM policy.
- B. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Directory. Use AWS SSO to configure an Active Directory user delegated to access the databases in RDS for SQL Server.
- C. Use Active Directory Connector to redirect directory requests to the company's on-premises Active Directory without caching any information in the cloud. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.
- D. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Directory. Ensure RDS for SQL Server is using mixed mode authentication. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.

**Answer: B**

**Explanation:**

[https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms\\_ad\\_tutorial\\_setup\\_trust.html](https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_tutorial_setup_trust.html)

#### **QUESTION 180**

A company uses an Amazon Redshift cluster to run its analytical workloads. Corporate policy requires that the company's data be encrypted at rest with customer managed keys. The company's disaster recovery plan requires that backups of the cluster be copied into another AWS Region on a regular basis.

How should a database specialist automate the process of backing up the cluster data in compliance with these policies?

- A. Copy the AWS Key Management Service (AWS KMS) customer managed key from the source Region to the destination Region. Set up an AWS Glue job in the source Region to copy the latest snapshot of the Amazon Redshift cluster from the source Region to the destination Region. Use a time-based schedule in AWS Glue to run the job on a daily basis.
- B. Create a new AWS Key Management Service (AWS KMS) customer managed key in the destination Region. Create a snapshot copy grant in the destination Region specifying the new key. In the source Region, configure cross-Region snapshots for the Amazon Redshift cluster specifying the destination Region, the snapshot copy grant, and retention periods for the snapshot.
- C. Copy the AWS Key Management Service (AWS KMS) customer-managed key from the source Region to the destination Region. Create Amazon S3 buckets in each Region using the keys from their respective Regions. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function in the source Region to copy the latest snapshot to the S3 bucket in that Region. Configure S3 Cross-Region Replication to copy the snapshots to the destination Region, specifying the source and destination KMS key IDs in the replication configuration.
- D. Use the same customer-supplied key materials to create a CMK with the same private key in the destination Region. Configure cross-Region snapshots in the source Region targeting the destination Region. Specify the corresponding CMK in the destination Region to encrypt the



snapshot.

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/kms/latest/developerguide/kms-dg.pdf>

#### QUESTION 181

A database specialist is launching a test graph database using Amazon Neptune for the first time. The database specialist needs to insert millions of rows of test observations from a .csv file that is stored in Amazon S3. The database specialist has been using a series of API calls to upload the data to the Neptune DB instance. Which combination of steps would allow the database specialist to upload the data faster? (Choose 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 database's loader endpoint.

**Answer:** BCF

**Explanation:**

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

<https://docs.aws.amazon.com/neptune/latest/userguide/dms-neptune.html>

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html#bulk-load-prereqs-s3>

#### QUESTION 182

A company is using Amazon DynamoDB global tables for an online gaming application. The game has players around the world. As the game has become more popular, the volume of requests to DynamoDB has increased significantly. Recently, players have reported that the game state is inconsistent between players in different countries. A database specialist observes that the ReplicationLatency metric for some of the replica tables is too high. Which approach will alleviate the problem?

- A. Configure all replica tables to use DynamoDB auto scaling.
- B. Configure a DynamoDB Accelerator (DAX) cluster on each of the replicas.
- C. Configure the primary table to use DynamoDB auto scaling and the replica tables to use manually provisioned capacity.
- D. Configure the table-level write throughput limit service quota to a higher value.

**Answer:** A

**Explanation:**

Using DynamoDB auto scaling is the recommended way to manage throughput capacity settings for replica tables that use the provisioned mode.

Reference:

[https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables\\_reqs\\_bestpractices.html](https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_reqs_bestpractices.html)

#### QUESTION 183

A company runs a MySQL database for its ecommerce application on a single Amazon RDS DB instance. Application purchases are automatically saved to the database, which causes intensive writes. Company employees frequently generate purchase reports. The company needs to improve database performance and reduce downtime due to patching for upgrades.

Which approach will meet these requirements with the LEAST amount of operational overhead?

- A. Enable a Multi-AZ deployment of the RDS for MySQL DB instance, and enable Memcached in the

[DBS-C01 Exam Dumps](#) [DBS-C01 Exam Questions](#) [DBS-C01 PDF Dumps](#) [DBS-C01 VCE Dumps](#)

<https://www.braindump2go.com/dbs-c01.html>

MySQL option group.

- B. Enable a Multi-AZ deployment of the RDS for MySQL DB instance, and set up replication to a MySQL DB instance running on Amazon EC2.
- C. Enable a Multi-AZ deployment of the RDS for MySQL DB instance, and add a read replica.
- D. Add a read replica and promote it to an Amazon Aurora MySQL DB cluster master. Then enable Amazon Aurora Serverless.

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

#### **QUESTION 184**

An ecommerce company is migrating its core application database to Amazon Aurora MySQL. The company is currently performing online transaction processing (OLTP) stress testing with concurrent database sessions. During the first round of tests, a database specialist noticed slow performance for some specific write operations. Reviewing Amazon CloudWatch metrics for the Aurora DB cluster showed 90% CPU utilization. Which steps should the database specialist take to MOST effectively identify the root cause of high CPU utilization and slow performance? (Choose two.)

- A. Enable Enhanced Monitoring at less than 30 seconds of granularity to review the operating system metrics before the next round of tests.
- B. Review the VolumeBytesUsed metric in CloudWatch to see if there is a spike in write I/O.
- C. Review Amazon RDS Performance Insights to identify the top SQL statements and wait events.
- D. Review Amazon RDS API calls in AWS CloudTrail to identify long-running queries.
- E. Enable Advance Auditing to log QUERY events in Amazon CloudWatch before the next round of tests.

**Answer: BC**

**Explanation:**

[https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_PerfInsights.API.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.API.html)

#### **QUESTION 185**

An online advertising company is implementing an application that displays advertisements to its users. The application uses an Amazon DynamoDB table as a data store. The application also uses a DynamoDB Accelerator (DAX) cluster to cache its reads. Most of the reads are from the GetItem query and the BatchGetItem query. Consistency of reads is not a requirement for this application. Upon deployment, the application cache is not performing as expected. Specific strongly consistent queries that run against the DAX cluster are taking many milliseconds to respond instead of microseconds. How can the company improve the cache behavior to increase application performance?

- A. Increase the size of the DAX cluster.
- B. Configure DAX to be an item cache with no query cache
- C. Use eventually consistent reads instead of strongly consistent reads.
- D. Create a new DAX cluster with a higher TTL for the item cache.

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadConsistency.html>

#### **QUESTION 186**

A company is running its critical production workload on a 500 GB Amazon Aurora MySQL DB cluster. A database engineer must move the workload to a new Amazon Aurora Serverless MySQL DB cluster without data loss. Which solution will accomplish the move with the LEAST downtime and the LEAST application impact?

- A. Modify the existing DB cluster and update the Aurora configuration to "Serverless."
- B. Create a snapshot of the existing DB cluster and restore it to a new Aurora Serverless DB cluster.

**[DBS-C01 Exam Dumps](#) [DBS-C01 Exam Questions](#) [DBS-C01 PDF Dumps](#) [DBS-C01 VCE Dumps](#)**

**<https://www.braindump2go.com/dbs-c01.html>**

- C. Create an Aurora Serverless replica from the existing DB cluster and promote it to primary when the replica lag is minimal.
- D. Replicate the data between the existing DB cluster and a new Aurora Serverless DB cluster by using AWS Database Migration Service (AWS DMS) with change data capture (CDC) enabled.

**Answer: C**

**Explanation:**

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

#### **QUESTION 187**

A company is building a web application on AWS. The application requires the database to support read and write operations in multiple AWS Regions simultaneously. The database also needs to propagate data changes between Regions as the changes occur. The application must be highly available and must provide latency of single-digit milliseconds.

Which solution meets these requirements?

- A. Amazon DynamoDB global tables
- B. Amazon DynamoDB streams with AWS Lambda to replicate the data
- C. An Amazon ElastiCache for Redis cluster with cluster mode enabled and multiple shards
- D. An Amazon Aurora global database

**Answer: A**

**Explanation:**

Global tables enable you to read and write your data locally providing single-digit-millisecond latency for your globally distributed application at any scale.

Reference: <https://aws.amazon.com/dynamodb/global-tables/>

#### **QUESTION 188**

A company is using Amazon Neptune as the graph database for one of its products. The company's data science team accidentally created large amounts of temporary information during an ETL process. The Neptune DB cluster automatically increased the storage space to accommodate the new data, but the data science team deleted the unused information.

What should a database specialist do to avoid unnecessary charges for the unused cluster volume space?

- A. Take a snapshot of the cluster volume. Restore the snapshot in another cluster with a smaller volume size.
- B. Use the AWS CLI to turn on automatic resizing of the cluster volume.
- C. Export the cluster data into a new Neptune DB cluster.
- D. Add a Neptune read replica to the cluster. Promote this replica as a new primary DB instance. Reset the storage space of the cluster.

**Answer: B**

**Explanation:**

In addition, the post offers programmatic approaches for automatically stopping or detecting idle resources that are incurring costs, allowing you to avoid unnecessary charges.

Reference: <https://aws.amazon.com/blogs/machine-learning/right-sizing-resources-and-avoiding-unnecessary-costs-in-amazon-sagemaker/>

#### **QUESTION 189**

A database specialist is responsible for designing a highly available solution for online transaction processing (OLTP) using Amazon RDS for MySQL production databases. Disaster recovery requirements include a cross-Region deployment along with an RPO of 5 minutes and RTO of 30 minutes.

What should the database specialist do to align to the high availability and disaster recovery requirements?

- A. Use a Multi-AZ deployment in each Region.
- B. Use read replica deployments in all Availability Zones of the secondary Region.



- C. Use Multi-AZ and read replica deployments within a Region.
- D. Use Multi-AZ and deploy a read replica in a secondary Region.

**Answer: C**

**Explanation:**

DR for Multi-AZ with in-Region read replicas – While Amazon RDS Multi-AZ provides HA and data protection, the associated in-Region read replica renders the scalability of read-only workloads, and the cross-Region automated backups feature provides DR.

Reference: <https://dataintegration.info/managed-disaster-recovery-with-amazon-rds-for-oracle-cross-region-automated-backups-part-1>

#### **QUESTION 190**

A media company wants to use zero-downtime patching (ZDP) for its Amazon Aurora MySQL database. Multiple processing applications are using SSL certificates to connect to database endpoints and the read replicas. Which factor will have the LEAST impact on the success of ZDP?

- A. Binary logging is enabled, or binary log replication is in progress.
- B. Current SSL connections are open to the database.
- C. Temporary tables or table locks are in use.
- D. The value of the `lower_case_table_names` server parameter was set to 0 when the tables were created.

**Answer: A**

**Explanation:**

Aurora MySQL 2.10 and higher, Aurora can perform a zero-downtime patch when binary log replication is enabled.

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Updates.Patching.html>

#### **QUESTION 191**

A financial services company has an application deployed on AWS that uses an Amazon Aurora PostgreSQL DB cluster. A recent audit showed that no log files contained database administrator activity. A database specialist needs to recommend a solution to provide database access and activity logs. The solution should use the least amount of effort and have a minimal impact on performance.

Which solution should the database specialist recommend?

- A. Enable Aurora Database Activity Streams on the database in synchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Kinesis Data Firehose destination to an Amazon S3 bucket.
- B. Create an AWS CloudTrail trail in the Region where the database runs. Associate the database activity logs with the trail.
- C. Enable Aurora Database Activity Streams on the database in asynchronous mode. Connect the Amazon Kinesis data stream to Kinesis Data Firehose. Set the Firehose destination to an Amazon S3 bucket.
- D. Allow connections to the DB cluster through a bastion host only. Restrict database access to the bastion host and application servers. Push the bastion host logs to Amazon CloudWatch Logs using the CloudWatch Logs agent.

**Answer: C**

**Explanation:**

<https://aws.amazon.com/blogs/database/filter-amazon-aurora-database-activity-stream-data-for-segregation-and-monitoring/>

#### **QUESTION 192**

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.

**[DBS-C01 Exam Dumps](#) [DBS-C01 Exam Questions](#) [DBS-C01 PDF Dumps](#) [DBS-C01 VCE Dumps](#)**

**<https://www.braindump2go.com/dbs-c01.html>**

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:**

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

#### **QUESTION 193**

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:**

**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>

#### **QUESTION 194**

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:**

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

#### **QUESTION 195**

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

**[DBS-C01 Exam Dumps](#) [DBS-C01 Exam Questions](#) [DBS-C01 PDF Dumps](#) [DBS-C01 VCE Dumps](#)**

**<https://www.braindump2go.com/dbs-c01.html>**

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:**

**Explanation:**

<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 196**

A financial services company runs an on-premises MySQL database for a critical application. The company is dissatisfied with its current database disaster recovery (DR) solution. The application experiences a significant amount of downtime whenever the database fails over to its DR facility. The application also experiences slower response times when reports are processed on the same database. To minimize the downtime in DR situations, the company has decided to migrate the database to AWS. The company requires a solution that is highly available and the most cost-effective.

Which solution meets these requirements?

- A. Create an Amazon RDS for MySQL Multi-AZ DB instance and configure a read replica in a different Availability Zone. Configure the application to reference the replica instance endpoint and report queries to reference the primary DB instance endpoint.
- B. Create an Amazon RDS for MySQL Multi-AZ DB instance and configure a read replica in a different Availability Zone. Configure the application to reference the primary DB instance endpoint and report queries to reference the replica instance endpoint.
- C. Create an Amazon Aurora DB cluster and configure an Aurora Replica in a different Availability Zone. Configure the application to reference the cluster endpoint and report queries to reference the reader endpoint.
- D. Create an Amazon Aurora DB cluster and configure an Aurora Replica in a different Availability Zone. Configure the application to reference the primary DB instance endpoint and report queries to reference the replica instance endpoint.

**Answer:** A

**Explanation:**

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

#### **QUESTION 197**

A company with 500,000 employees needs to supply its employee list to an application used by human resources. Every 30 minutes, the data is exported using the LDAP service to load into a new Amazon DynamoDB table. The data model has a base table with Employee ID for the partition key and a global secondary index with Organization ID as the partition key.

While importing the data, a database specialist receives ProvisionedThroughputExceededException errors. After increasing the provisioned write capacity units (WCUs) to 50,000, the specialist receives the same errors. Amazon CloudWatch metrics show a consumption of 1,500 WCUs.

What should the database specialist do to address the issue?

- A. Change the data model to avoid hot partitions in the global secondary index.
- B. Enable auto scaling for the table to automatically increase write capacity during bulk imports.
- C. Modify the table to use on-demand capacity instead of provisioned capacity.
- D. Increase the number of retries on the bulk loading application.

**Answer:** B

**Explanation:**

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

**QUESTION 198**

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:**

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

**QUESTION 199**

A company wants to build a new invoicing service for its cloud-native application on AWS. The company has a small development team and wants to focus on service feature development and minimize operations and maintenance as much as possible. The company expects the service to handle billions of requests and millions of new records every day. The service feature requirements, including data access patterns are well-defined. The service has an availability target of 99.99% with a milliseconds latency requirement. The database for the service will be the system of record for invoicing data.

Which database solution meets these requirements at the LOWEST cost?

- A. Amazon Neptune
- B. Amazon Aurora PostgreSQL Serverless
- C. Amazon RDS for PostgreSQL
- D. Amazon DynamoDB

**Answer:** A

**Explanation:**

<https://aws.amazon.com/neptune>

**QUESTION 200**

Application developers have reported that an application is running slower as more users are added. The application database is running on an Amazon Aurora DB cluster with an Aurora Replica. The application is written to take advantage of read scaling through reader endpoints. A database specialist looks at the performance metrics of the database and determines that, as new users were added to the database, the primary instance CPU utilization steadily increased while the Aurora Replica CPU utilization remained steady.

How can the database specialist improve database performance while ensuring minimal downtime?

- A. Modify the Aurora DB cluster to add more replicas until the overall load stabilizes. Then, reduce the number of replicas once the application meets service level objectives.
- B. Modify the primary instance to a larger instance size that offers more CPU capacity.
- C. Modify a replica to a larger instance size that has more CPU capacity. Then, promote the modified replica.
- D. Restore the Aurora DB cluster to one that has an instance size with more CPU capacity. Then, swap the names of the old and new DB clusters.

**[DBS-C01 Exam Dumps](#) [DBS-C01 Exam Questions](#) [DBS-C01 PDF Dumps](#) [DBS-C01 VCE Dumps](#)**

**<https://www.braindump2go.com/dbs-c01.html>**

**Answer:** B

**Explanation:**

The automatic storage increase setting of a primary instance automatically applies to any read replicas of that instance. The automatic storage increase setting cannot be independently set for read replicas.

Reference: <https://cloud.google.com/sql/docs/postgres/instance-settings>

#### **QUESTION 201**

A company's development team needs to have production data restored in a staging AWS account. The production database is running on an Amazon RDS for PostgreSQL Multi-AZ DB instance, which has AWS KMS encryption enabled using the default KMS key. A database specialist planned to share the most recent automated snapshot with the staging account, but discovered that the option to share snapshots is disabled in the AWS Management Console. What should the database specialist do to resolve this?

- A. Disable automated backups in the DB instance. Share both the automated snapshot and the default KMS key with the staging account. Restore the snapshot in the staging account and enable automated backups.
- B. Copy the automated snapshot specifying a custom KMS encryption key. Share both the copied snapshot and the custom KMS encryption key with the staging account. Restore the snapshot to the staging account within the same Region.
- C. Modify the DB instance to use a custom KMS encryption key. Share both the automated snapshot and the custom KMS encryption key with the staging account. Restore the snapshot in the staging account.
- D. Copy the automated snapshot while keeping the default KMS key. Share both the snapshot and the default KMS key with the staging account. Restore the snapshot in the staging account.

**Answer:** B

**Explanation:**

You then create a copy of that snapshot and specify a KMS key to encrypt that snapshot copy. You can then restore an encrypted DB instance from the encrypted snapshot.

Reference: [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_CopySnapshot.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CopySnapshot.html)

#### **QUESTION 202**

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.

**Answer:** D

**Explanation:**

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

#### **QUESTION 203**

A retail company uses Amazon Redshift Spectrum to run complex analytical queries on objects that are stored in an Amazon S3 bucket. The objects are joined with multiple dimension tables that are stored in an Amazon Redshift database. The company uses the database to create monthly and quarterly aggregated reports. Users who attempt to run queries are reporting the following error message:

**[DBS-C01 Exam Dumps](#) **[DBS-C01 Exam Questions](#) **[DBS-C01 PDF Dumps](#) **[DBS-C01 VCE Dumps](#)********

**<https://www.braindump2go.com/dbs-c01.html>**



error: Spectrum Scan Error: Access throttled  
Which solution will resolve this error?

- A. Check file sizes of fact tables in Amazon S3, and look for large files. Break up large files into smaller files of equal size between 100 MB and 1 GB
- B. Reduce the number of queries that users can run in parallel.
- C. Check file sizes of fact tables in Amazon S3, and look for small files. Merge the small files into larger files of at least 64 MB in size.
- D. Review and optimize queries that submit a large aggregation step to Redshift Spectrum.

**Answer: A**

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files>

#### **QUESTION 204**

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 205**

An application reads and writes data to an Amazon RDS for MySQL DB instance. A new reporting dashboard needs read-only access to the database. When the application and reports are both under heavy load, the database experiences performance degradation. A database specialist needs to improve the database performance.

What should the database specialist do to meet these requirements?

- A. Create a read replica of the DB instance. Configure the reports to connect to the replication instance endpoint.
- B. Create a read replica of the DB instance. Configure the application and reports to connect to the cluster endpoint.
- C. Enable Multi-AZ deployment. Configure the reports to connect to the standby replica.
- D. Enable Multi-AZ deployment. Configure the application and reports to connect to the cluster endpoint.

**Answer: C**

**Explanation:**

<https://www.nops.io/enable-multi-az-for-rds-instances/>