

**AWS Certified Database - Specialty.VCEplus.premium.exam.85q**

Passing Score: 800  
Time Limit: 120 min  
File Version: 1.0



**Website:** <https://vceplus.com>

**VCE to PDF Converter:** <https://vceplus.com/vce-to-pdf/>

**Facebook:** <https://www.facebook.com/VCE.For.All.VN/>

**Twitter :** [https://twitter.com/VCE\\_Plus](https://twitter.com/VCE_Plus)

**AWS Certified Database – Specialty (beta)**

**Version 1.0**



## Exam A

### QUESTION 1

A company has deployed an e-commerce web application in a new AWS account. An Amazon RDS for MySQL Multi-AZ DB instance is part of this deployment with a database-1.xxxxxxxxxx.us-east-1.rds.amazonaws.com endpoint listening on port 3306. The company's Database Specialist is able to log in to MySQL and run queries from the bastion host using these details.

When users try to utilize the application hosted in the AWS account, they are presented with a generic error message. The application servers are logging a "could not connect to server: Connection times out" error message to Amazon CloudWatch Logs.

What is the cause of this error?

- A. The user name and password the application is using are incorrect.
- B. The security group assigned to the application servers does not have the necessary rules to allow inbound connections from the DB instance.
- C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.
- D. The user name and password are correct, but the user is not authorized to use the DB instance.

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Reference: <https://forums.aws.amazon.com/thread.jspa?threadID=129700>

### QUESTION 2

An AWS CloudFormation stack that included an Amazon RDS DB instance was accidentally deleted and recent data was lost. A Database Specialist needs to add RDS settings to the CloudFormation template to reduce the chance of accidental instance data loss in the future.

Which settings will meet this requirement? (Choose 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



**Correct Answer: ACF**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html> <https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-accidental-updates/>

### QUESTION 3

A Database Specialist is troubleshooting an application connection failure on an Amazon Aurora DB cluster with multiple Aurora Replicas that had been running with no issues for the past 2 months. The connection failure lasted for 5 minutes and corrected itself after that. The Database Specialist reviewed the Amazon RDS events and determined a failover event occurred at that time. The failover process took around 15 seconds to complete.

What is the MOST likely cause of the 5-minute connection outage?

- A. After a database crash, Aurora needed to replay the redo log from the last database checkpoint
- B. The client-side application is caching the DNS data and its TTL is set too high
- C. After failover, the Aurora DB cluster needs time to warm up before accepting client connections
- D. There were no active Aurora Replicas in the Aurora DB cluster

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

**QUESTION 4**

A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being set over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine.

Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster. Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses.

What should the Database Specialist do to correct the Data Analysts' inability to connect?

- A. Restart the DB cluster to apply the SSL change.
- B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.
- C. Add explicit mappings between the Data Analysts' IP addresses and the instance in the security group assigned to the DB cluster.
- D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 5**

A company is concerned about the cost of a large-scale, transactional application using Amazon DynamoDB that only needs to store data for 2 days before it is deleted. In looking at the tables, a Database Specialist notices that much of the data is months old, and goes back to when the application was first deployed.

What can the Database Specialist do to reduce the overall cost?

- A. Create a new attribute in each table to track the expiration time and create an AWS Glue transformation to delete entries more than 2 days old.
- B. Create a new attribute in each table to track the expiration time and enable DynamoDB Streams on each table.
- C. Create a new attribute in each table to track the expiration time and enable time to live (TTL) on each table.
- D. Create an Amazon CloudWatch Events event to export the data to Amazon S3 daily using AWS Data Pipeline and then truncate the Amazon DynamoDB table.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 6**

A company has an on-premises system that tracks various database operations that occur over the lifetime of a database, including database shutdown, deletion, creation, and backup.

The company recently moved two databases to Amazon RDS and is looking at a solution that would satisfy these requirements. The data could be used by other systems within the company.

Which solution will meet these requirements with minimal effort?

- A. Create an Amazon Cloudwatch Events rule with the operations that need to be tracked on Amazon RDS. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.
- B. Create an AWS Lambda function to trigger on AWS CloudTrail API calls. Filter on specific RDS API calls and write the output to the tracking systems.
- C. Create RDS event subscriptions. Have the tracking systems subscribe to specific RDS event system notifications.
- D. Write RDS logs to Amazon Kinesis Data Firehose. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 7**

A clothing company uses a custom ecommerce application and a PostgreSQL database to sell clothes to thousands of users from multiple countries. The company is migrating its application and database from its on-premises data center to the AWS Cloud. The company has selected Amazon EC2 for the application and Amazon RDS for PostgreSQL for the database. The company requires database passwords to be changed every 60 days. A Database Specialist needs to ensure that the credentials used by the web application to connect to the database are managed securely.

Which approach should the Database Specialist take to securely manage the database credentials?

- A. Store the credentials in a text file in an Amazon S3 bucket. Restrict permissions on the bucket to the IAM role associated with the instance profile only. Modify the application to download the text file and retrieve the credentials on startup. Update the text file every 60 days.
- B. Configure IAM database authentication for the application to connect to the database. Create an IAM user and map it to a separate database user for each ecommerce user. Require users to update their passwords every 60 days.
- C. Store the credentials in AWS Secrets Manager. Restrict permissions on the secret to only the IAM role associated with the instance profile. Modify the application to retrieve the credentials from Secrets Manager on start up. Configure the rotation interval to 60 days.
- D. Store the credentials in an encrypted text file in the application AMI. Use AWS KMS to store the key for decrypting the text file. Modify the application to decrypt the text file and retrieve the credentials on start up. Update the text file and publish a new AMI every 60 days.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 8

A financial services company is developing a shared data service that supports different applications from throughout the company. A Database Specialist designed a solution to leverage Amazon ElastiCache for Redis with cluster mode enabled to enhance performance and scalability. The cluster is configured to listen on port 6379.

Which combination of steps should the Database Specialist take to secure the cache data and protect it from unauthorized access? (Choose three.)

- A. Enable in-transit and at-rest encryption on the ElastiCache cluster.
- B. Ensure that Amazon CloudWatch metrics are configured in the ElastiCache cluster.
- C. Ensure the security group for the ElastiCache cluster allows all inbound traffic from itself and inbound traffic on TCP port 6379 from trusted clients only.
- D. Create an IAM policy to allow the application service roles to access all ElastiCache API actions.
- E. Ensure the security group for the ElastiCache clients authorize inbound TCP port 6379 and port 22 traffic from the trusted ElastiCache cluster's security group.
- F. Ensure the cluster is created with the auth-token parameter and that the parameter is used in all subsequent commands.

**Correct Answer:** ABE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/getting-started/tutorials/setting-up-a-redis-cluster-with-amazon-elasticache/>

**QUESTION 9** A company is running an Amazon RDS for PostgreSQL DB instance and wants to migrate it to an Amazon Aurora PostgreSQL DB cluster. The current database is 1 TB in size. The migration needs to have minimal downtime.

What is the FASTEST way to accomplish this?

- A. Create an Aurora PostgreSQL DB cluster. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.
- B. Use the pg\_dump and pg\_restore utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.
- C. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.
- D. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replica. Promote the replica during the cutover.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 10**

A Database Specialist is migrating a 2 TB Amazon RDS for Oracle DB instance to an RDS for PostgreSQL DB instance using AWS DMS. The source RDS Oracle DB instance is in a VPC in the us-east-1 Region. The target RDS for PostgreSQL DB instance is in a VPC in the use-west-2 Region.

Where should the AWS DMS replication instance be placed for the MOST optimal performance?

- A. In the same Region and VPC of the source DB instance
- B. In the same Region and VPC as the target DB instance
- C. In the same VPC and Availability Zone as the target DB instance
- D. In the same VPC and Availability Zone as the source DB instance

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 11

The Development team recently executed a database script containing several data definition language (DDL) and data manipulation language (DML) statements on an Amazon Aurora MySQL DB cluster. The release accidentally deleted thousands of rows from an important table and broke some application functionality. This was discovered 4 hours after the release. Upon investigation, a Database Specialist tracked the issue to a DELETE command in the script with an incorrect WHERE clause filtering the wrong set of rows.

The Aurora DB cluster has Backtrack enabled with an 8-hour backtrack window. The Database Administrator also took a manual snapshot of the DB cluster before the release started. The database needs to be returned to the correct state as quickly as possible to resume full application functionality. Data loss must be minimal. How can the Database Specialist accomplish this?

- A. Quickly rewind the DB cluster to a point in time before the release using Backtrack.
- B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database.
- C. Restore the DB cluster using the manual backup snapshot created before the release and change the application configuration settings to point to the new DB cluster.
- D. Create a clone of the DB cluster with Backtrack enabled. Rewind the cloned cluster to a point in time before the release. Copy deleted rows from the clone to the original database.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 12

A company is load testing its three-tier production web application deployed with an AWS CloudFormation template on AWS. The Application team is making changes to deploy additional Amazon EC2 and AWS Lambda resources to expand the load testing capacity. A Database Specialist wants to ensure that the changes made by the Application team will not change the Amazon RDS database resources already deployed.

Which combination of steps would allow the Database Specialist to accomplish this? (Choose two.)

- A. Review the stack drift before modifying the template
- B. Create and review a change set before applying it
- C. Export the database resources as stack outputs
- D. Define the database resources in a nested stack
- E. Set a stack policy for the database resources

**Correct Answer:** AD

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 13

A manufacturing company's website uses an Amazon Aurora PostgreSQL DB cluster.

Which configurations will result in the LEAST application downtime during a failover? (Choose 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.

**Correct Answer:** ABC

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 14

A company is hosting critical business data in an Amazon Redshift cluster. Due to the sensitive nature of the data, the cluster is encrypted at rest using AWS KMS. As a part of disaster recovery requirements, the company needs to copy the Amazon Redshift snapshots to another Region.

Which steps should be taken in the AWS Management Console to meet the disaster recovery requirements?

- A. Create a new KMS customer master key in the source Region. Switch to the destination Region, enable Amazon Redshift cross-Region snapshots, and use the KMS key of the source Region.
- B. Create a new IAM role with access to the KMS key. Enable Amazon Redshift cross-Region replication using the new IAM role, and use the KMS key of the source Region.
- C. Enable Amazon Redshift cross-Region snapshots in the source Region, and create a snapshot copy grant and use a KMS key in the destination Region.
- D. Create a new KMS customer master key in the destination Region and create a new IAM role with access to the new KMS key. Enable Amazon Redshift cross-Region replication in the source Region and use the KMS key of the destination Region.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-snapshots.html>



#### QUESTION 15

A company has a production Amazon Aurora DB cluster that serves both online transaction processing (OLTP) transactions and compute-intensive reports. The reports run for 10% of the total cluster uptime while the OLTP transactions run all the time. The company has benchmarked its workload and determined that a six-node Aurora DB cluster is appropriate for the peak workload.

The company is now looking at cutting costs for this DB cluster, but needs to have a sufficient number of nodes in the cluster to support the workload at different times. The workload has not changed since the previous benchmarking exercise.

How can a Database Specialist address these requirements with minimal user involvement?

- A. Split up the DB cluster into two different clusters: one for OLTP and the other for reporting. Monitor and set up replication between the two clusters to keep data consistent.
- B. Review and evaluate the peak combined workload. Ensure that utilization of the DB cluster node is at an acceptable level. Adjust the number of instances, if necessary.
- C. Use the stop cluster functionality to stop all the nodes of the DB cluster during times of minimal workload. The cluster can be restarted again depending on the workload at the time.
- D. Set up automatic scaling on the DB cluster. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 16

A company is running a finance application on an Amazon RDS for MySQL DB instance. The application is governed by multiple financial regulatory agencies. The RDS DB instance is set up with security groups to allow access to certain Amazon EC2 servers only. AWS KMS is used for encryption at rest.

Which step will provide additional security?

- A. Set up NACLs that allow the entire EC2 subnet to access the DB instance

- B. Disable the master user account
- C. Set up a security group that blocks SSH to the DB instance
- D. Set up RDS to use SSL for data in transit

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/database/applying-best-practices-for-securing-sensitive-data-in-amazon-rds/>

#### QUESTION 17

A company needs a data warehouse solution that keeps data in a consistent, highly structured format. The company requires fast responses for end-user queries when looking at data from the current year, and users must have access to the full 15-year dataset, when needed. This solution also needs to handle a fluctuating number incoming queries. Storage costs for the 100 TB of data must be kept low.

Which solution meets these requirements?

- A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storage. Provision enough instances to support high demand.
- B. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Provision enough instances to support high demand.
- C. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Enable Amazon Redshift Concurrency Scaling.
- D. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Leverage Amazon Redshift elastic resize.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**



#### QUESTION 18

A gaming company wants to deploy a game in multiple Regions. The company plans to save local high scores in Amazon DynamoDB tables in each Region. A Database Specialist needs to design a solution to automate the deployment of the database with identical configurations in additional Regions, as needed. The solution should also automate configuration changes across all Regions.

Which solution would meet these requirements and deploy the DynamoDB tables?

- A. Create an AWS CLI command to deploy the DynamoDB table to all the Regions and save it for future deployments.
- B. Create an AWS CloudFormation template and deploy the template to all the Regions.
- C. Create an AWS CloudFormation template and use a stack set to deploy the template to all the Regions.
- D. Create DynamoDB tables using the AWS Management Console in all the Regions and create a step-by-step guide for future deployments.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 19

A team of Database Specialists is currently investigating performance issues on an Amazon RDS for MySQL DB instance and is reviewing related metrics. The team wants to narrow the possibilities down to specific database wait events to better understand the situation.

How can the Database Specialists accomplish this?

- A. Enable the option to push all database logs to Amazon CloudWatch for advanced analysis
- B. Create appropriate Amazon CloudWatch dashboards to contain specific periods of time
- C. Enable Amazon RDS Performance Insights and review the appropriate dashboard
- D. Enable Enhanced Monitoring with the appropriate settings

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 20

A large company is using an Amazon RDS for Oracle Multi-AZ DB instance with a Java application. As a part of its disaster recovery annual testing, the company would like to simulate an Availability Zone failure and record how the application reacts during the DB instance failover activity. The company does not want to make any code changes for this activity.

What should the company do to achieve this in the shortest amount of time?

- A. Use a blue-green deployment with a complete application-level failover test
- B. Use the RDS console to reboot the DB instance by choosing the option to reboot with failover
- C. Use RDS fault injection queries to simulate the primary node failure
- D. Add a rule to the NACL to deny all traffic on the subnets associated with a single Availability Zone

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: [https://wellarchitectedlabs.com/Reliability/300\\_Testing\\_for\\_Resiliency\\_of\\_EC2\\_RDS\\_and\\_S3/Lab\\_Guide.html](https://wellarchitectedlabs.com/Reliability/300_Testing_for_Resiliency_of_EC2_RDS_and_S3/Lab_Guide.html)

#### QUESTION 21

A company maintains several databases using Amazon RDS for MySQL and PostgreSQL. Each RDS database generates log files with retention periods set to their default values. The company has now mandated that database logs be maintained for up to 90 days in a centralized repository to facilitate real-time and after-the-fact analyses.

What should a Database Specialist do to meet these requirements with minimal effort?

- A. Create an AWS Lambda function to pull logs from the RDS databases and consolidate the log files in an Amazon S3 bucket. Set a lifecycle policy to expire the objects after 90 days.
- B. Modify the RDS databases to publish log to Amazon CloudWatch Logs. Change the log retention policy for each log group to expire the events after 90 days.
- C. Write a stored procedure in each RDS database to download the logs and consolidate the log files in an Amazon S3 bucket. Set a lifecycle policy to expire the objects after 90 days.
- D. Create an AWS Lambda function to download the logs from the RDS databases and publish the logs to Amazon CloudWatch Logs. Change the log retention policy for the log group to expire the events after 90 days.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 22

A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium-sized primary instance, one large-sized replica, and two medium sized replicas. The Database Specialist did not assign a promotion tier to the replicas.

In the event of a primary failure, what will occur?

- A. Aurora will promote an Aurora Replica that is of the same size as the primary instance
- B. Aurora will promote an arbitrary Aurora Replica
- C. Aurora will promote the largest-sized Aurora Replica
- D. Aurora will not promote an Aurora Replica

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-ug.pdf>

**QUESTION 23** A company is running its line of business application on AWS, which uses Amazon RDS for MySQL at the persistent data store. The company wants to minimize downtime when it migrates the database to Amazon Aurora.

Which migration method should a Database Specialist use?

- A. Take a snapshot of the RDS for MySQL DB instance and create a new Aurora DB cluster with the option to migrate snapshots.
- B. Make a backup of the RDS for MySQL DB instance using the mysqldump utility, create a new Aurora DB cluster, and restore the backup.
- C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.
- D. Create a clone of the RDS for MySQL DB instance and promote the Aurora DB cluster.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://d1.awsstatic.com/whitepapers/RDS/Migrating%20your%20databases%20to%20Amazon%20Aurora.pdf> (10)

**QUESTION 24**

The Security team for a finance company was notified of an internal security breach that happened 3 weeks ago. A Database Specialist must start producing audit logs out of the production Amazon Aurora PostgreSQL cluster for the Security team to use for monitoring and alerting. The Security team is required to perform real-time alerting and monitoring outside the Aurora DB cluster and wants to have the cluster push encrypted files to the chosen solution. Which approach will meet these requirements?

- A. Use pg\_audit to generate audit logs and send the logs to the Security team.
- B. Use AWS CloudTrail to audit the DB cluster and the Security team will get data from Amazon S3.
- C. Set up database activity streams and connect the data stream from Amazon Kinesis to consumer applications.
- D. Turn on verbose logging and set up a schedule for the logs to be dumped out for the Security team.

**Correct Answer:** B

**Section:** (none)

**Explanation**



**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-ug.pdf> (525)

**QUESTION 25**

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 to 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

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 26**

A media company is using Amazon RDS for PostgreSQL to store user data. The RDS DB instance currently has a publicly accessible setting enabled and is hosted in a public subnet. Following a recent AWS Well-Architected Framework review, a Database Specialist was given new security requirements.

- Only certain on-premises corporate network IPs should connect to the DB instance. ▪

Connectivity is allowed from the corporate network only.

Which combination of steps does the Database Specialist need to take to meet these new requirements? (Choose three.)

- A. Modify the pg\_hba.conf file. Add the required corporate network IPs and remove the unwanted IPs.
- B. Modify the associated security group. Add the required corporate network IPs and remove the unwanted IPs.
- C. Move the DB instance to a private subnet using AWS DMS.
- D. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
- E. Disable the publicly accessible setting.
- F. Connect to the DB instance using private IPs and a VPN.

**Correct Answer:** DEF

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 27

A company is about to launch a new product, and test databases must be re-created from production data. The company runs its production databases on an Amazon Aurora MySQL DB cluster. A Database Specialist needs to deploy a solution to create these test databases as quickly as possible with the least amount of administrative effort.

What should the Database Specialist do to meet these requirements?

- A. Restore a snapshot from the production cluster into test clusters
- B. Create logical dumps of the production cluster and restore them into new test clusters
- C. Use database cloning to create clones of the production cluster
- D. Add an additional read replica to the production cluster and use that node for testing

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**



#### QUESTION 28

A company with branch offices in Portland, New York, and Singapore has a three-tier web application that leverages a shared database. The database runs on Amazon RDS for MySQL and is hosted in the us-west-2 Region. The application has a distributed front end deployed in the us-west-2, ap-southeast-1, and us-east-2 Regions.

This front end is used as a dashboard for Sales Managers in each branch office to see current sales statistics. There are complaints that the dashboard performs more slowly in the Singapore location than it does in Portland or New York. A solution is needed to provide consistent performance for all users in each location.

Which set of actions will meet these requirements?

- A. Take a snapshot of the instance in the us-west-2 Region. Create a new instance from the snapshot in the ap-southeast-1 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- B. Create an RDS read replica in the ap-southeast-1 Region from the primary RDS DB instance in the us-west-2 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- C. Create a new RDS instance in the ap-southeast-1 Region. Use AWS DMS and change data capture (CDC) to update the new instance in the ap-southeast-1 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- D. Create an RDS read replica in the us-west-2 Region where the primary instance resides. Create a read replica in the ap-southeast-1 Region from the read replica located on the us-west-2 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 29**

A company wants to migrate its existing on-premises Oracle database to Amazon Aurora PostgreSQL. The migration must be completed with minimal downtime using AWS DMS. A Database Specialist must validate that the data was migrated accurately from the source to the target before the cutover. The migration must have minimal impact on the performance of the source database. Which approach will MOST effectively meet these requirements?

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert source Oracle database schemas to the target Aurora DB cluster. Verify the datatype of the columns.
- B. Use the table metrics of the AWS DMS task created for migrating the data to verify the statistics for the tables being migrated and to verify that the data definition language (DDL) statements are completed.
- C. Enable the AWS Schema Conversion Tool (AWS SCT) premigration validation and review the premigration checklist to make sure there are no issues with the conversion.
- D. Enable AWS DMS data validation on the task so the AWS DMS task compares the source and target records, and reports any mismatches.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: [https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Validating.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Validating.html)

**QUESTION 30**

A company is planning to close for several days. A Database Specialist needs to stop all applications along with the DB instances to ensure employees do not have access to the systems during this time. All databases are running on Amazon RDS for MySQL.

The Database Specialist wrote and executed a script to stop all the DB instances. When reviewing the logs, the Database Specialist found that Amazon RDS DB instances with read replicas did not stop.

How should the Database Specialist edit the script to fix this issue?

- A. Stop the source instances before stopping their read replicas
- B. Delete each read replica before stopping its corresponding source instance
- C. Stop the read replicas before stopping their source instances
- D. Use the AWS CLI to stop each read replica and source instance at the same

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 31**

A global digital advertising company captures browsing metadata to contextually display relevant images, pages, and links to targeted users. A single page load can generate multiple events that need to be stored individually. The maximum size of an event is 200 KB and the average size is 10 KB. Each page load must query the user's browsing history to provide targeting recommendations. The advertising company expects over 1 billion page visits per day from users in the United States, Europe, Hong Kong, and India. The structure of the metadata varies depending on the event. Additionally, the browsing metadata must be written and read with very low latency to ensure a good viewing experience for the users.

Which database solution meets these requirements?

- A. Amazon DocumentDB
- B. Amazon RDS Multi-AZ deployment
- C. Amazon DynamoDB global table
- D. Amazon Aurora Global Database

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

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

**QUESTION 32**

A Database Specialist modified an existing parameter group currently associated with a production Amazon RDS for SQL Server Multi-AZ DB instance. The change is associated with a static parameter type, which controls the number of user connections allowed on the most critical RDS SQL Server DB instance for the company. This change has been approved for a specific maintenance window to help minimize the impact on users.

How should the Database Specialist apply the parameter group change for the DB instance?

- A. Select the option to apply the change immediately
- B. Allow the preconfigured RDS maintenance window for the given DB instance to control when the change is applied
- C. Apply the change manually by rebooting the DB instance during the approved maintenance window
- D. Reboot the secondary Multi-AZ DB instance

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: [https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Troubleshooting.html](https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Troubleshooting.html)

#### QUESTION 33

A Database Specialist is designing a new database infrastructure for a ride hailing application. The application data includes a ride tracking system that stores GPS coordinates for all rides. Real-time statistics and metadata lookups must be performed with high throughput and microsecond latency. The database should be fault tolerant with minimal operational overhead and development effort. Which solution meets these requirements in the MOST efficient way?

- A. Use Amazon RDS for MySQL as the database and use Amazon ElastiCache
- B. Use Amazon DynamoDB as the database and use DynamoDB Accelerator
- C. Use Amazon Aurora MySQL as the database and use Aurora's buffer cache
- D. Use Amazon DynamoDB as the database and use Amazon API Gateway

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/solutions/case-studies/lyft/>

#### QUESTION 34

A company is using an Amazon Aurora PostgreSQL DB cluster with an xlarge primary instance master and two large Aurora Replicas for high availability and read-only workload scaling. A failover event occurs and application performance is poor for several minutes. During this time, application servers in all Availability Zones are healthy and responding normally.

What should the company do to eliminate this application performance issue?

- A. Configure both of the Aurora Replicas to the same instance class as the primary DB instance. Enable cache coherence on the DB cluster, set the primary DB instance failover priority to tier-0, and assign a failover priority of tier-1 to the replicas.
- B. Deploy an AWS Lambda function that calls the DescribeDBInstances action to establish which instance has failed, and then use the PromoteReadReplica operation to promote one Aurora Replica to be the primary DB instance. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.
- C. Configure one Aurora Replica to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance class. Set the failover priority to tier-1 for the other replicas.
- D. Configure both Aurora Replicas to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 35

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.

**Correct Answer:** BE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 36

A company is using Amazon with Aurora Replicas for read-only workload scaling. A Database Specialist needs to split up two read-only applications so each application always connects to a dedicated replica. The Database Specialist wants to implement load balancing and high availability for the read-only applications. Which solution meets these requirements?

- A. Use a specific instance endpoint for each replica and add the instance endpoint to each read-only application connection string.
- B. Use reader endpoints for both the read-only workload applications.
- C. Use a reader endpoint for one read-only application and use an instance endpoint for the other read-only application.
- D. Use custom endpoints for the two read-only applications.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://rimzy.net/category/amazon-rds/page/4/>

#### QUESTION 37

An online gaming company is planning to launch a new game with Amazon DynamoDB as its data store. The database should be designated to support the following use cases:

- Update scores in real time whenever a player is playing the game.
- Retrieve a player's score details for a specific game session.

A Database Specialist decides to implement a DynamoDB table. Each player has a unique `user_id` and each game has a unique `game_id`.

Which choice of keys is recommended for the DynamoDB table?

- A. Create a global secondary index with `game_id` as the partition key
- B. Create a global secondary index with `user_id` as the partition key
- C. Create a composite primary key with `game_id` as the partition key and `user_id` as the sort key
- D. Create a composite primary key with `user_id` as the partition key and `game_id` as the sort key

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/database/amazon-dynamodb-gaming-use-cases-and-design-patterns/>

#### QUESTION 38

A Database Specialist migrated an existing production MySQL database from on-premises to an Amazon RDS for MySQL DB instance. However, after the migration, the database needed to be encrypted at rest using AWS KMS. Due to the size of the database, reloading the data into an encrypted database would be too time-consuming, so it is not an option. How should the Database Specialist satisfy this new requirement?

- A. Create a snapshot of the unencrypted RDS DB instance. Create an encrypted copy of the unencrypted snapshot. Restore the encrypted snapshot copy.
- B. Modify the RDS DB instance. Enable the AWS KMS encryption option that leverages the AWS CLI.
- C. Restore an unencrypted snapshot into a MySQL RDS DB instance that is encrypted.
- D. Create an encrypted read replica of the RDS DB instance. Promote it the master.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 39**

A Database Specialist is planning to create a read replica of an existing Amazon RDS for MySQL Multi-AZ DB instance. When using the AWS Management Console to conduct this task, the Database Specialist discovers that the source RDS DB instance does not appear in the read replica source selection box, so the read replica cannot be created. What is the most likely reason for this?

- A. The source DB instance has to be converted to Single-AZ first to create a read replica from it.
- B. Enhanced Monitoring is not enabled on the source DB instance.
- C. The minor MySQL version in the source DB instance does not support read replicas.
- D. Automated backups are not enabled on the source DB instance.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/rds/features/read-replicas/>

**QUESTION 40**

A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL. The schema and the data have been migrated successfully. The on-premises database server was also being used to run database maintenance cron jobs written in Python to perform tasks including data purging and generating data exports. The logs for these jobs show that, most of the time, the jobs completed within 5 minutes, but a few jobs took up to 10 minutes to complete. These maintenance jobs need to be set up for Aurora PostgreSQL.

How can the Database Specialist schedule these jobs so the setup requires minimal maintenance and provides high availability?

- A. Create cron jobs on an Amazon EC2 instance to run the maintenance jobs following the required schedule.
- B. Connect to the Aurora host and create cron jobs to run the maintenance jobs following the required schedule.
- C. Create AWS Lambda functions to run the maintenance jobs and schedule them with Amazon CloudWatch Events.
- D. Create the maintenance job using the Amazon CloudWatch job scheduling plugin.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/systems-manager/latest/userguide/mw-cli-task-options.html>

**QUESTION 41**

A company has an Amazon RDS Multi-AZ DB instances that is 200 GB in size with an RPO of 6 hours. To meet the company's disaster recovery policies, the database backup needs to be copied into another Region. The company requires the solution to be cost-effective and operationally efficient.

What should a Database Specialist do to copy the database backup into a different Region?

- A. Use Amazon RDS automated snapshots and use AWS Lambda to copy the snapshot into another Region
- B. Use Amazon RDS automated snapshots every 6 hours and use Amazon S3 cross-Region replication to copy the snapshot into another Region
- C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a second Lambda function to copy the snapshot into another Region
- D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot of the read replica

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/>

**QUESTION 42**

An Amazon RDS EBS-optimized instance with Provisioned IOPS (PIOPS) storage is using less than half of its allocated IOPS over the course of several hours under constant load. The RDS instance exhibits multi-second read and write latency, and uses all of its maximum bandwidth for read throughput, yet the instance uses less than half of its CPU and RAM resources. What should a Database Specialist do in this situation to increase performance and return latency to sub-second levels?

- A. Increase the size of the DB instance storage



- B. Change the underlying EBS storage type to General Purpose SSD (gp2)
- C. Disable EBS optimization on the DB instance
- D. Change the DB instance to an instance class with a higher maximum bandwidth

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 43

After restoring an Amazon RDS snapshot from 3 days ago, a company's Development team cannot connect to the restored RDS DB instance. What is the likely cause of this problem?

- A. The restored DB instance does not have Enhanced Monitoring enabled
- B. The production DB instance is using a custom parameter group
- C. The restored DB instance is using the default security group
- D. The production DB instance is using a custom option group

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 44

A gaming company has implemented a leaderboard in AWS using a Sorted Set data structure within Amazon ElastiCache for Redis. The ElastiCache cluster has been deployed with cluster mode disabled and has a replication group deployed with two additional replicas. The company is planning for a worldwide gaming event and is anticipating a higher write load than what the current cluster can handle.

Which method should a Database Specialist use to scale the ElastiCache cluster ahead of the upcoming event?

- A. Enable cluster mode on the existing ElastiCache cluster and configure separate shards for the Sorted Set across all nodes in the cluster.
- B. Increase the size of the ElastiCache cluster nodes to a larger instance size.
- C. Create an additional ElastiCache cluster and load-balance traffic between the two clusters.
- D. Use the EXPIRE command and set a higher time to live (TTL) after each call to increment a given key.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/database/work-with-cluster-mode-on-amazon-elasticache-for-redis/>

#### QUESTION 45

An ecommerce company has tasked a Database Specialist with creating a reporting dashboard that visualizes critical business metrics that will be pulled from the core production database running on Amazon Aurora. Data that is read by the dashboard should be available within 100 milliseconds of an update.

The Database Specialist needs to review the current configuration of the Aurora DB cluster and develop a cost-effective solution. The solution needs to accommodate the unpredictable read workload from the reporting dashboard without any impact on the write availability and performance of the DB cluster. Which solution meets these requirements?

- A. Turn on the serverless option in the DB cluster so it can automatically scale based on demand.
- B. Provision a clone of the existing DB cluster for the new Application team.
- C. Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).
- D. Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:****QUESTION 46**

A retail company is about to migrate its online and mobile store to AWS. The company's CEO has strategic plans to grow the brand globally. A Database Specialist has been challenged to provide predictable read and write database performance with minimal operational overhead.

What should the Database Specialist do to meet these requirements?

- A. Use Amazon DynamoDB global tables to synchronize transactions
- B. Use Amazon EMR to copy the orders table data across Regions
- C. Use Amazon Aurora Global Database to synchronize all transactions
- D. Use Amazon DynamoDB Streams to replicate all DynamoDB transactions and sync them

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

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

**QUESTION 47**

A company is closing one of its remote data centers. This site runs a 100 TB on-premises data warehouse solution. The company plans to use the AWS Schema Conversion Tool (AWS SCT) and AWS DMS for the migration to AWS. The site network bandwidth is 500 Mbps. A Database Specialist wants to migrate the on-premises data using Amazon S3 as the data lake and Amazon Redshift as the data warehouse. This move must take place during a 2-week period when source systems are shut down for maintenance. The data should stay encrypted at rest and in transit. Which approach has the least risk and the highest likelihood of a successful data transfer?

- A. Set up a VPN tunnel for encrypting data over the network from the data center to AWS. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Once complete, start an AWS DMS task to move the data from the source to Amazon S3. Use AWS Glue to load the data from Amazon S3 to Amazon Redshift.
- B. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryption. Use AWS DMS to finish copying data to Amazon Redshift.
- C. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Once complete, use a fleet of 10 TB dedicated encrypted drives using the AWS Import/Export feature to copy data from on-premises to Amazon S3 with AWS KMS encryption. Use AWS Glue to load the data to Amazon Redshift.
- D. Set up a VPN tunnel for encrypting data over the network from the data center to AWS. Leverage a native database export feature to export the data and compress the files. Use the aws S3 cp multi-port upload command to upload these files to Amazon S3 with AWS KMS encryption. Once complete, load the data to Amazon Redshift using AWS Glue.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:****QUESTION 48**

A company is looking to migrate a 1 TB Oracle database from on-premises to an Amazon Aurora PostgreSQL DB cluster. The company's Database Specialist discovered that the Oracle database is storing 100 GB of large binary objects (LOBs) across multiple tables. The Oracle database has a maximum LOB size of 500 MB with an average LOB size of 350 MB. The Database Specialist has chosen AWS DMS to migrate the data with the largest replication instances. How should the Database Specialist optimize the database migration using AWS DMS?

- A. Create a single task using full LOB mode with a LOB chunk size of 500 MB to migrate the data and LOBs together
- B. Create two tasks: task1 with LOB tables using full LOB mode with a LOB chunk size of 500 MB and task2 without LOBs
- C. Create two tasks: task1 with LOB tables using limited LOB mode with a maximum LOB size of 500 MB and task 2 without LOBs
- D. Create a single task using limited LOB mode with a maximum LOB size of 500 MB to migrate data and LOBs together

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 49**

A Database Specialist is designing a disaster recovery strategy for a production Amazon DynamoDB table. The table uses provisioned read/write capacity mode, global secondary indexes, and time to live (TTL). The Database Specialist has restored the latest backup to a new table.

To prepare the new table with identical settings, which steps should be performed? (Choose two.)

- A. Re-create global secondary indexes in the new table
- B. Define IAM policies for access to the new table
- C. Define the TTL settings
- D. Encrypt the table from the AWS Management Console or use the update-table command
- E. Set the provisioned read and write capacity

**Correct Answer:** AE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

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

**QUESTION 50**

A Database Specialist is creating Amazon DynamoDB tables, Amazon CloudWatch alarms, and associated infrastructure for an Application team using a development AWS account. The team wants a deployment method that will standardize the core solution components while managing environment-specific settings separately, and wants to minimize rework due to configuration errors. Which process should the Database Specialist recommend to meet these requirements?

- A. Organize common and environmental-specific parameters hierarchically in the AWS Systems Manager Parameter Store, then reference the parameters dynamically from an AWS CloudFormation template. Deploy the CloudFormation stack using the environment name as a parameter.
- B. Create a parameterized AWS CloudFormation template that builds the required objects. Keep separate environment parameter files in separate Amazon S3 buckets. Provide an AWS CLI command that deploys the CloudFormation stack directly referencing the appropriate parameter bucket.
- C. Create a parameterized AWS CloudFormation template that builds the required objects. Import the template into the CloudFormation interface in the AWS Management Console. Make the required changes to the parameters and deploy the CloudFormation stack.
- D. Create an AWS Lambda function that builds the required objects using an AWS SDK. Set the required parameter values in a test event in the Lambda console for each environment that the Application team can modify, as needed. Deploy the infrastructure by triggering the test event in the console.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/mt/aws-cloudformation-signed-sealed-and-deployed/>

**QUESTION 51**

A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi-AZ DB instance. Tests were run on the database after work hours, which generated additional database logs. The free storage of the RDS DB instance is low due to these additional logs.

What should the company do to address this space constraint issue?

- A. Log in to the host and run the `rm $PGDATA/pg_logs/*` command
- B. Modify the `rds.log_retention_period` parameter to 1440 and wait up to 24 hours for database logs to be deleted
- C. Create a ticket with AWS Support to have the logs deleted
- D. Run the `SELECT rds_rotate_error_log()` stored procedure to rotate the logs

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 52**

A user has a non-relational key-value database. The user is looking for a fully managed AWS service that will offload the administrative burdens of operating and scaling distributed databases. The solution must be cost-effective and able to handle unpredictable application traffic.

What should a Database Specialist recommend for this user?

- A. Create an Amazon DynamoDB table with provisioned capacity mode
- B. Create an Amazon DocumentDB cluster
- C. Create an Amazon DynamoDB table with on-demand capacity mode
- D. Create an Amazon Aurora Serverless DB cluster

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

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

#### QUESTION 53

A gaming company is designing a mobile gaming app that will be accessed by many users across the globe. The company wants to have replication and full support for multi-master writes. The company also wants to ensure low latency and consistent performance for app users. Which solution meets these requirements?

- A. Use Amazon DynamoDB global tables for storage and enable DynamoDB automatic scaling
- B. Use Amazon Aurora for storage and enable cross-Region Aurora Replicas
- C. Use Amazon Aurora for storage and cache the user content with Amazon ElastiCacheD. Use Amazon Neptune for storage

**Correct Answer: A**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/database/how-to-use-amazon-dynamodb-global-tables-to-power-multiregion-architectures/>

#### QUESTION 54

A Database Specialist needs to speed up any failover that might occur on an Amazon Aurora PostgreSQL DB cluster. The Aurora DB cluster currently includes the primary instance and three Aurora Replicas. How can the Database Specialist ensure that failovers occur with the least amount of downtime for the application?

- A. Set the TCP keepalive parameters low
- B. Call the AWS CLI failover-db-cluster command
- C. Enable Enhanced Monitoring on the DB cluster
- D. Start a database activity stream on the DB cluster

**Correct Answer: B**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

#### QUESTION 55

A Database Specialist needs to define a database migration strategy to migrate an on-premises Oracle database to an Amazon Aurora MySQL DB cluster. The company requires near-zero downtime for the data migration. The solution must also be cost-effective.

Which approach should the Database Specialist take?

- A. Dump all the tables from the Oracle database into an Amazon S3 bucket using datapump (expdp). Run data transformations in AWS Glue. Load the data from the S3 bucket to the Aurora DB cluster.
- B. Order an AWS Snowball appliance and copy the Oracle backup to the Snowball appliance. Once the Snowball data is delivered to Amazon S3, create a new Aurora DB cluster. Enable the S3 integration to migrate the data directly from Amazon S3 to Amazon RDS.
- C. Use the AWS Schema Conversion Tool (AWS SCT) to help rewrite database objects to MySQL during the schema migration. Use AWS DMS to perform the full load and change data capture (CDC) tasks.
- D. Use AWS Server Migration Service (AWS SMS) to import the Oracle virtual machine image as an Amazon EC2 instance. Use the Oracle Logical Dump utility to migrate the Oracle data from Amazon EC2 to an Aurora DB cluster.

**Correct Answer: D**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

**QUESTION 56**

A marketing company is using Amazon DocumentDB and requires that database audit logs be enabled. A Database Specialist needs to configure monitoring so that all data definition language (DDL) statements performed are visible to the Administrator. The Database Specialist has set the audit\_logs parameter to enabled in the cluster parameter group.

What should the Database Specialist do to automatically collect the database logs for the Administrator?

- A. Enable DocumentDB to export the logs to Amazon CloudWatch Logs
- B. Enable DocumentDB to export the logs to AWS CloudTrail
- C. Enable DocumentDB Events to export the logs to Amazon CloudWatch Logs
- D. Configure an AWS Lambda function to download the logs using the download-db-log-file-portion operation and store the logs in Amazon S3

**Correct Answer: A**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/documentdb/latest/developerguide/profiling.html>

**QUESTION 57**

A company is looking to move an on-premises IBM Db2 database running AIX on an IBM POWER7 server. Due to escalating support and maintenance costs, the company is exploring the option of moving the workload to an Amazon Aurora PostgreSQL DB cluster.

What is the quickest way for the company to gather data on the migration compatibility?

- A. Perform a logical dump from the Db2 database and restore it to an Aurora DB cluster. Identify the gaps and compatibility of the objects migrated by comparing row counts from source and target tables.
- B. Run AWS DMS from the Db2 database to an Aurora DB cluster. Identify the gaps and compatibility of the objects migrated by comparing the row counts from source and target tables.
- C. Run native PostgreSQL logical replication from the Db2 database to an Aurora DB cluster to evaluate the migration compatibility.
- D. Run the AWS Schema Conversion Tool (AWS SCT) from the Db2 database to an Aurora DB cluster. Create a migration assessment report to evaluate the migration compatibility.

**Correct Answer: D**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/Schema-Conversion-Tool.pdf>

**QUESTION 58**

An ecommerce company is using Amazon DynamoDB as the backend for its order-processing application. The steady increase in the number of orders is resulting in increased DynamoDB costs. Order verification and reporting perform many repeated GetItem functions that pull similar datasets, and this read activity is contributing to the increased costs. The company wants to control these costs without significant development efforts. How should a Database Specialist address these requirements?

- A. Use AWS DMS to migrate data from DynamoDB to Amazon DocumentDB
- B. Use Amazon DynamoDB Streams and Amazon Kinesis Data Firehose to push the data into Amazon Redshift
- C. Use an Amazon ElastiCache for Redis in front of DynamoDB to boost read performance. Use DynamoDB Accelerator to offload the reads

**Correct Answer: B**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

**QUESTION 59**

An IT consulting company wants to reduce costs when operating its development environment databases. The company's workflow creates multiple Amazon Aurora MySQL DB clusters for each development group. The Aurora DB clusters are only used for 8 hours a day. The DB clusters can then be deleted at the end of the development cycle, which lasts 2 weeks. Which of the following provides the MOST cost-effective solution?

- A. Use AWS CloudFormation templates. Deploy a stack with the DB cluster for each development group. Delete the stack at the end of the development cycle.
- B. Use the Aurora DB cloning feature. Deploy a single development and test Aurora DB instance, and create clone instances for the development groups. Delete the clones at the end of the development cycle.
- C. Use Aurora Replicas. From the master automatic pause compute capacity option, create replicas for each development group, and promote each replica to master. Delete the replicas at the end of the development cycle.
- D. Use Aurora Serverless. Restore current Aurora snapshot and deploy to a serverless cluster for each development group. Enable the option to pause the compute capacity on the cluster and set an appropriate timeout.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 60

A company has multiple applications serving data from a secure on-premises database. The company is migrating all applications and databases to the AWS Cloud. The IT Risk and Compliance department requires that auditing be enabled on all secure databases to capture all log ins, log outs, failed logins, permission changes, and database schema changes. A Database Specialist has recommended Amazon Aurora MySQL as the migration target, and leveraging the Advanced Auditing feature in Aurora.

Which events need to be specified in the Advanced Auditing configuration to satisfy the minimum auditing requirements? (Choose three.)

- A. CONNECT
- B. QUERY\_DCL
- C. QUERY\_DDL
- D. QUERY\_DML
- E. TABLE
- F. QUERY

**Correct Answer:** ACE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

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

#### QUESTION 61

A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season.

Which solution will meet these requirements at the lowest cost?

- A. DynamoDB Streams
- B. DynamoDB with DynamoDB Accelerator
- C. DynamoDB with on-demand capacity mode
- D. DynamoDB with provisioned capacity mode with Auto Scaling

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: [https://aws.amazon.com/blogs/database/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/?nc1=b\\_rp](https://aws.amazon.com/blogs/database/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/?nc1=b_rp)

#### QUESTION 62

A company's Security department established new requirements that state internal users must connect to an existing Amazon RDS for SQL Server DB instance using their corporate Active Directory (AD) credentials. A Database Specialist must make the modifications needed to fulfill this requirement.

Which combination of actions should the Database Specialist take? (Choose three.)

- A. Disable Transparent Data Encryption (TDE) on the RDS SQL Server DB instance.



- B. Modify the RDS SQL Server DB instance to use the directory for Windows authentication. Create appropriate new logins.
- C. Use the AWS Management Console to create an AWS Managed Microsoft AD. Create a trust relationship with the corporate AD.
- D. Stop the RDS SQL Server DB instance, modify it to use the directory for Windows authentication, and start it again. Create appropriate new logins.
- E. Use the AWS Management Console to create an AD Connector. Create a trust relationship with the corporate AD.
- F. Configure the AWS Managed Microsoft AD domain controller Security Group.

**Correct Answer:** CDF

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 63

A Database Specialist is performing a proof of concept with Amazon Aurora using a small instance to confirm a simple database behavior. When loading a large dataset and creating the index, the Database Specialist encounters the following error message from Aurora:

```
ERROR: cloud not write block 7507718 of temporary file: No space left on device
```

What is the cause of this error and what should the Database Specialist do to resolve this issue?

- A. The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to modify the workload to load the data slowly.
- B. The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to enable Aurora storage scaling.
- C. The local storage used to store temporary tables is full. The Database Specialist needs to scale up the instance.
- D. The local storage used to store temporary tables is full. The Database Specialist needs to enable local storage scaling.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://serverfault.com/questions/109828/how-can-i-tune-postgres-to-avoid-this-error>



#### QUESTION 64

A financial company wants to store sensitive user data in an Amazon Aurora PostgreSQL DB cluster. The database will be accessed by multiple applications across the company. The company has mandated that all communications to the database be encrypted and the server identity must be validated. Any non-SSL-based connections should be disallowed access to the database.

Which solution addresses these requirements?

- A. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=allow`.
- B. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=disable`.
- C. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-ca`.
- D. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-full`.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://forums.aws.amazon.com/message.jspa?messageID=734076>

#### QUESTION 65

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:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 66

A company wants to automate the creation of secure test databases with random credentials to be stored safely for later use. The credentials should have sufficient information about each test database to initiate a connection and perform automated credential rotations. The credentials should not be logged or stored anywhere in an unencrypted form. Which steps should a Database Specialist take to meet these requirements using an AWS CloudFormation template?

- A. Create the database with the MasterUserName and MasterUserPassword properties set to the default values. Then, create the secret with the user name and password set to the same default values. Add a Secret Target Attachment resource with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database. Finally, update the secret's password value with a randomly generated string set by the GenerateSecretString property.
- B. Add a Mapping property from the database Amazon Resource Name (ARN) to the secret ARN. Then, create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property. Add the database with the MasterUserName and MasterUserPassword properties set to the user name of the secret.
- C. Add a resource of type AWS::SecretsManager::Secret and specify the GenerateSecretString property. Then, define the database user name in the SecureStringTemplate template. Create a resource for the database and reference the secret string for the MasterUserName and MasterUserPassword properties. Then, add a resource of type AWS::SecretsManager::SecretTargetAttachment with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database.
- D. Create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property. Add an SecretTargetAttachment resource with the SecretId property set to the Amazon Resource Name (ARN) of the secret and the TargetId property set to a parameter value matching the desired database ARN. Then, create a database with the MasterUserName and MasterUserPassword properties set to the previously created values in the secret.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-secrettargetattachment.html>

#### QUESTION 67

A company is going to use an Amazon Aurora PostgreSQL DB cluster for an application backend. The DB cluster contains some tables with sensitive data. A Database Specialist needs to control the access privileges at the table level. How can the Database Specialist meet these requirements?

- A. Use AWS IAM database authentication and restrict access to the tables using an IAM policy.
- B. Configure the rules in a NACL to restrict outbound traffic from the Aurora DB cluster.
- C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.
- D. Define access privileges to the tables containing sensitive data in the pg\_hba.conf file.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/database/managing-postgresql-users-and-roles/>

#### QUESTION 68

A Database Specialist is working with a company to launch a new website built on Amazon Aurora with several Aurora Replicas. This new website will replace an on-premises website connected to a legacy relational database. Due to stability issues in the legacy database, the company would like to test the resiliency of Aurora. Which action can the Database Specialist take to test the resiliency of the Aurora DB cluster?

- A. Stop the DB cluster and analyze how the website responds
- B. Use Aurora fault injection to crash the master DB instance
- C. Remove the DB cluster endpoint to simulate a master DB instance failure
- D. Use Aurora Backtrack to crash the DB cluster

**Correct Answer:** B

**Section: (none)**

**Explanation**

**Explanation/Reference:**

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

**QUESTION 69**

A company just migrated to Amazon Aurora PostgreSQL from an on-premises Oracle database. After the migration, the company discovered there is a period of time every day around 3:00 PM where the response time of the application is noticeably slower. The company has narrowed down the cause of this issue to the database and not the application. Which set of steps should the Database Specialist take to most efficiently find the problematic PostgreSQL query?

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption. Watch these dashboards during the next slow period.
- B. Launch an Amazon EC2 instance, and install and configure an open-source PostgreSQL monitoring tool that will run reports based on the output error logs.
- C. Modify the logging database parameter to log all the queries related to locking in the database and then check the logs after the next slow period for this information.
- D. Enable Amazon RDS Performance Insights on the PostgreSQL database. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

**Correct Answer: D**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

**QUESTION 70**

A company has a web-based survey application that uses Amazon DynamoDB. During peak usage, when survey responses are being collected, a Database Specialist sees the ProvisionedThroughputExceededException error. What can the Database Specialist do to resolve this error? (Choose two.)

- A. Change the table to use Amazon DynamoDB Streams
- B. Purchase DynamoDB reserved capacity in the affected Region
- C. Increase the write capacity units for the specific table
- D. Change the table capacity mode to on-demand
- E. Change the table type to throughput optimized



**Correct Answer: CE**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Reference: <https://forums.aws.amazon.com/thread.jspa?threadID=174315>

**QUESTION 71**

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

**Correct Answer: ACD**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

**QUESTION 72**

A company is building a new web platform where user requests trigger an AWS Lambda function that performs an insert into an Amazon Aurora MySQL DB cluster. Initial tests with less than 10 users on the new platform yielded successful execution and fast response times. However, upon more extensive tests with the actual target of 3,000 concurrent users, Lambda functions are unable to connect to the DB cluster and receive too many connections errors. Which of the following will resolve this issue?

- A. Edit the my.cnf file for the DB cluster to increase max\_connections
- B. Increase the instance size of the DB cluster
- C. Change the DB cluster to Multi-AZ
- D. Increase the number of Aurora Replicas

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 73**

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

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 74**

A company is writing a new survey application to be used with a weekly televised game show. The application will be available for 2 hours each week. The company expects to receive over 500,000 entries every week, with each survey asking 2-3 multiple choice questions of each user. A Database Specialist needs to select a platform that is highly scalable for a large number of concurrent writes to handle the anticipated volume. Which AWS services should the Database Specialist consider? (Choose two.)

- A. Amazon DynamoDB
- B. Amazon Redshift
- C. Amazon Neptune
- D. Amazon Elasticsearch Service
- E. Amazon ElastiCache

**Correct Answer:** AE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 75**

A company has migrated a single MySQL database to Amazon Aurora. The production data is hosted in a DB cluster in VPC\_PROD, and 12 testing environments are hosted in VPC\_TEST using the same AWS account. Testing results in minimal changes to the test data. The Development team wants each environment refreshed nightly so each test database contains fresh production data every day.

Which migration approach will be the fastest and most cost-effective to implement?

- A. Run the master in Amazon Aurora MySQL. Create 12 clones in VPC\_TEST, and script the clones to be deleted and re-created nightly.

- B. Run the master in Amazon Aurora MySQL. Take a nightly snapshot, and restore it into 12 databases in VPC\_TEST using Aurora Serverless.
- C. Run the master in Amazon Aurora MySQL. Create 12 Aurora Replicas in VPC\_TEST, and script the replicas to be deleted and re-created nightly.
- D. Run the master in Amazon Aurora MySQL using Aurora Serverless. Create 12 clones in VPC\_TEST, and script the clones to be deleted and re-created nightly.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 76

A large ecommerce company uses Amazon DynamoDB to handle the transactions on its web portal. Traffic patterns throughout the year are usually stable; however, a large event is planned. The company knows that traffic will increase by up to 10 times the normal load over the 3-day event. When sale prices are published during the event, traffic will spike rapidly.

How should a Database Specialist ensure DynamoDB can handle the increased traffic?

- A. Ensure the table is always provisioned to meet peak needs
- B. Allow burst capacity to handle the additional load
- C. Set an AWS Application Auto Scaling policy for the table to handle the increase in traffic
- D. Preprovision additional capacity for the known peaks and then reduce the capacity after the event

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 77

A Database Specialist is migrating an on-premises Microsoft SQL Server application database to Amazon RDS for PostgreSQL using AWS DMS. The application requires minimal downtime when the RDS DB instance goes live. What change should the Database Specialist make to enable the migration?

- A. Configure the on-premises application database to act as a source for an AWS DMS full load with ongoing change data capture (CDC)
- B. Configure the AWS DMS replication instance to allow both full load and ongoing change data capture (CDC)
- C. Configure the AWS DMS task to generate full logs to allow for ongoing change data capture (CDC)
- D. Configure the AWS DMS connections to allow two-way communication to allow for ongoing change data capture (CDC)

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/rds-import-data/>

#### QUESTION 78

A financial company has allocated an Amazon RDS MariaDB DB instance with large storage capacity to accommodate migration efforts. Post-migration, the company purged unwanted data from the instance. The company now want to downsize storage to save money. The solution must have the least impact on production and near-zero downtime. Which solution would meet these requirements?

- A. Create a snapshot of the old databases and restore the snapshot with the required storage
- B. Create a new RDS DB instance with the required storage and move the databases from the old instances to the new instance using AWS DMSC. Create a new database using native backup and restore
- D. Create a new read replica and make it the primary by terminating the existing primary

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 79**

A large financial services company requires that all data be encrypted in transit. A Developer is attempting to connect to an Amazon RDS DB instance using the company VPC for the first time with credentials provided by a Database Specialist. Other members of the Development team can connect, but this user is consistently receiving an error indicating a communications link failure. The Developer asked the Database Specialist to reset the password a number of times, but the error persists.

Which step should be taken to troubleshoot this issue?

- A. Ensure that the database option group for the RDS DB instance allows ingress from the Developer machine's IP address
- B. Ensure that the RDS DB instance's subnet group includes a public subnet to allow the Developer to connect
- C. Ensure that the RDS DB instance has not reached its maximum connections limit
- D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 80**

A company is running Amazon RDS for MySQL for its workloads. There is downtime when AWS operating system patches are applied during the Amazon RDS-specified maintenance window.

What is the MOST cost-effective action that should be taken to avoid downtime?

- A. Migrate the workloads from Amazon RDS for MySQL to Amazon DynamoDB
- B. Enable cross-Region read replicas and direct read traffic to them when Amazon RDS is down
- C. Enable a read replicas and direct read traffic to it when Amazon RDS is down
- D. Enable an Amazon RDS for MySQL Multi-AZ configuration

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 81**

A Database Specialist must create a read replica to isolate read-only queries for an Amazon RDS for MySQL DB instance. Immediately after creating the read replica, users that query it report slow response times. What could be causing these slow response times?

- A. New volumes created from snapshots load lazily in the background
- B. Long-running statements on the master
- C. Insufficient resources on the master
- D. Overload of a single replication thread by excessive writes on the master

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 82**

A company developed an AWS CloudFormation template used to create all new Amazon DynamoDB tables in its AWS account. The template configures provisioned throughput capacity using hard-coded values. The company wants to change the template so that the tables it creates in the future have independently configurable read and write capacity units assigned. Which solution will enable this change?

- A. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Configure DynamoDB to provision throughput capacity using the stack's mappings.
- B. Add values for two Number parameters, rcuCount and wcuCount, to the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.
- C. Add values for the rcuCount and wcuCount parameters as outputs of the template. Configure DynamoDB to provision throughput capacity using the stack outputs.
- D. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.



**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 83

A retail company with its main office in New York and another office in Tokyo plans to build a database solution on AWS. The company's main workload consists of a mission-critical application that updates its application data in a data store. The team at the Tokyo office is building dashboards with complex analytical queries using the application data. The dashboards will be used to make buying decisions, so they need to have access to the application data in less than 1 second.

Which solution meets these requirements?

- A. Use an Amazon RDS DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region. Create an Amazon ElastiCache cluster in the ap-northeast-1 Region to cache application data from thereplica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in the us-east-1 Region with replication into the ap-northeast-1 Region. Use Amazon QuickSight for displaying dashboard results.
- C. Use an Amazon RDS for MySQL DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region. Have the dashboard application read from the read replica.
- D. Use an Amazon Aurora global database. Deploy the writer instance in the us-east-1 Region and the replica in the ap-northeast-1 Region. Have the dashboard application read from the replica ap-northeast-1 Region.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 84

A company is using Amazon RDS for PostgreSQL. The Security team wants all database connection requests to be logged and retained for 180 days. The RDS for PostgreSQL DB instance is currently using the default parameter group. A Database Specialist has identified that setting the log\_connections parameter to 1 will enable connections logging.

Which combination of steps should the Database Specialist take to meet the logging and retention requirements? (Choose two.)

- A. Update the log\_connections parameter in the default parameter group
- B. Create a custom parameter group, update the log\_connections parameter, and associate the parameter with the DB instance
- C. Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to 180 days
- D. Enable publishing of database engine logs to an Amazon S3 bucket and set the lifecycle policy to 180 days
- E. Connect to the RDS PostgreSQL host and update the log\_connections parameter in the postgresql.conf file

**Correct Answer:** AE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Reference: <https://aws.amazon.com/blogs/database/working-with-rds-and-aurora-postgresql-logs-part-1/>

#### QUESTION 85

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

**Correct Answer:** BD

**Section:** (none)

**Explanation**

**Explanation/Reference:**

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

