Latest Version: 8.0

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.xxxxxxxxxxxx.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.

Answer: C

Explanation:
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

Answer: ACF

Explanation:

Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attributedeletionpolicy.
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

Answer: B

Explanation:
When your application tries to establish a connection after a failover, the new Aurora PostgreSQL writer will be a previous reader, which can be found using the Aurora read only endpoint before DNS updates have fully propagated. Setting the java DNS TTL to a low value helps cycle between reader nodes on subsequent connection attempts.
Amazon Aurora is designed to recover from a crash almost instantaneously and continue to serve your application data. Unlike other databases, after a crash Amazon Aurora does not need to replay the redo log from the last database checkpoint before making the database available for operations. Amazon Aurora performs crash recovery asynchronously on parallel threads, so your database is open and available immediately after a crash. Because the storage is organized in many small segments, each with its own redo log, the underlying storage can replay redo records on demand in parallel and asynchronously as part of a disk read after a crash. This approach reduces database restart times to less than 60 seconds in most cases

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.

Answer: B

Explanation:
• To connect using SSL:
• Provide the SSLTrust certificate (can be downloaded from AWS)
• Provide SSL options when connecting to database
• Not using SSL on a DB that enforces SSL would result in error
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/ssl-certificate-rotation-aurorapostgresql.html

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.

Answer: C

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html