Q&A

Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt.
----- Guaranteed.

KILL EXAMS

PASS ✓

Amazon

SAA-C03

AWS Certified Solutions Architect - Associate

Question: 84

A Solutions Architect is building a cloud infrastructure where EC2 instances require access to various AWS services such as S3 and Redshift. The Architect will also need to provide access to system

administrators so they can deploy and test their changes.

Which configuration should be used to ensure that the access to the resources is secured and not compromised? (Select TWO.)
A. Store the AWS Access Keys in the EC2 instance.
B. Assign an IAM role to the Amazon EC2 instance.
C. Store the AWS Access Keys in AC
E. Enable Multi-Factor Authentication.
F. Assign an IAM user for each Amazon EC2 Instance.

Answer: B,D

Explanation:

In this scenario, the correct answers are:

– Enable Multi-Factor Authentication

– Assign an IAM role to the Amazon EC2 instance

Always remember that you should associate IAM roles to EC2 instances and not an IAM user, for the purpose of accessing other AWS services. IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles.
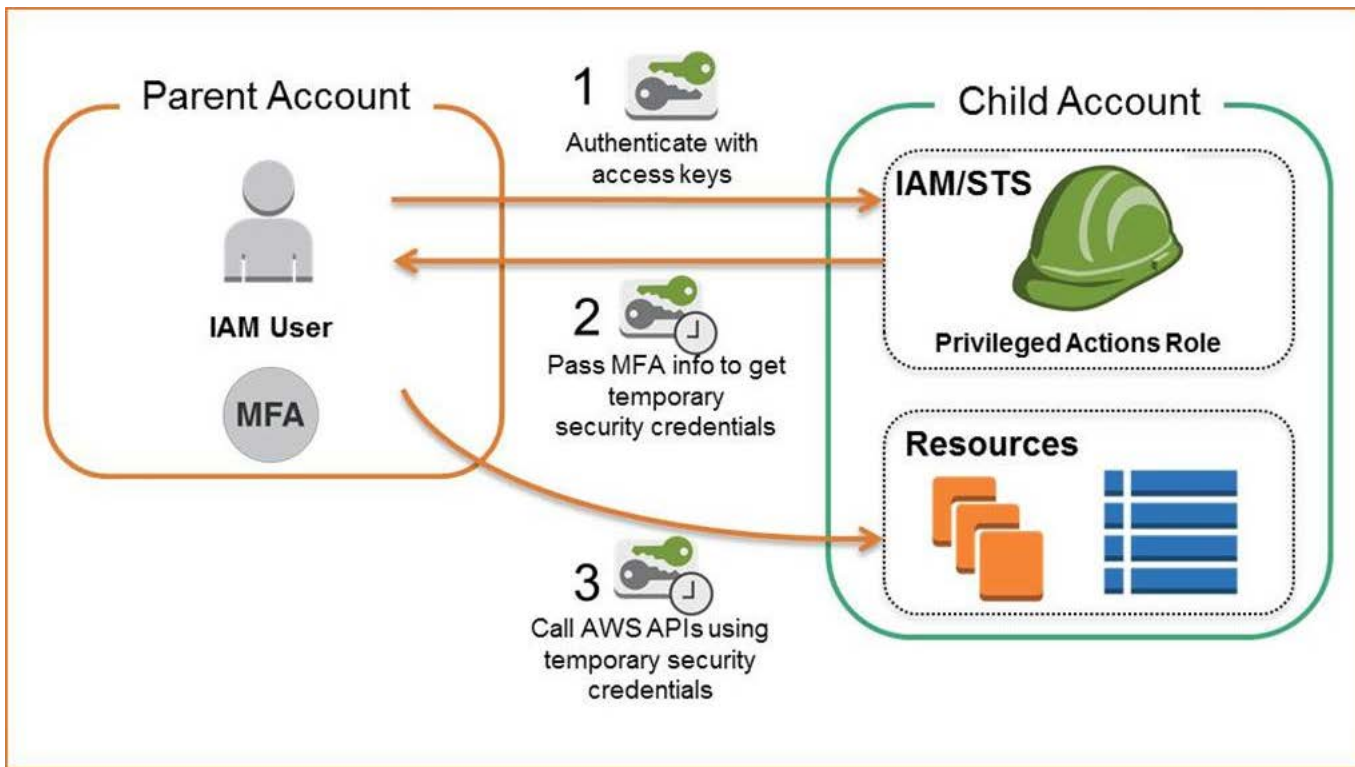
AWS Multi-Factor Authentication (MFA) is a simple best practice that adds an extra layer of protection on top of your user name and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for their user name and password (the first factor—what they know), as well as for an authentication code from their AWS MFA device (the second factor—what they have). Taken together, these multiple factors provide increased security for your AWS account settings and resources. You can enable MFA for your AWS account and for individual IAM users you have created under your account. MFA can also be used to control access to AWS service APIs.

Storing the AWS Access Keys in the EC2 instance is incorrect. This is not recommended by AWS as it can be compromised. Instead of storing access keys on an EC2 instance for use by applications that run on the instance and make AWS API requests, you can use an IAM role to provide temporary access keys for these applications.

Assigning an IAM user for each Amazon EC2 Instance is incorrect because there is no need to create an IAM user for this scenario since IAM roles already provide greater flexibility and easier management. Storing the AWS Access Keys in ACM is incorrect because ACM is just a service that lets you easily provision, manage, and deploy public and private SSL/TLS certificates for use with AWS services and your internal connected resources. It is not used as a secure storage for your access keys. References:

https://aws.amazon.com/iam/details/mfa/

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

Check out this AWS IAM Cheat Sheet:

https://tutorialsdojo.com/aws-identity-and-access-management-iam/

Question: 85

A company needs to deploy at least 2 EC2 instances to support the normal workloads of its application and automatically scale up to 6 EC2 instances to handle the peak load. The architecture must be highly available and fault-tolerant as it is processing mission-critical workloads.

As the Solutions Architect of the company, what should you do to meet the above requirement?

A. Create an Auto Scaling group of EC2 instances and set the minimum capacity to 2 and the maximum capacity to 4. Deploy 2 instances in Availability Zone A and 2 instances in Availability Zone B.

B. Create an Auto Scaling group of EC2 instances and set the minimum capacity to 4 and the maximum capacity to 6. Deploy 2 instances in Availability Zone A and another 2 instances in Availability Zone B.

C. Create an Auto Scaling group of EC2 instances and set the minimum capacity to 2 and the maximum capacity to 6. Deploy 4 instances in Availability Zone A.

D. Create an Auto Scaling group of EC2 instances and set the minimum capacity to 2 and the

maximum capacity to 6. Use 2 Availability Zones and deploy 1 instance for each AZ.

Answer: B

Explanation:

Amazon EC2 Auto Scaling helps ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can also specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size.
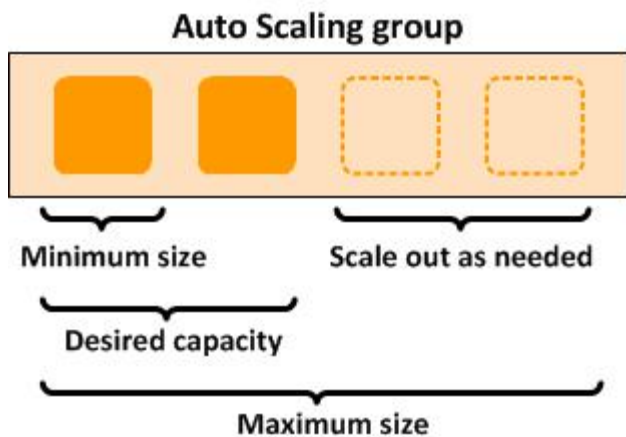
To achieve highly available and fault-tolerant architecture for your applications, you must deploy all your instances in different Availability Zones. This will help you isolate your resources if an outage occurs. Take note that to achieve fault tolerance, you need to have redundant resources in place to avoid any system degradation in the event of a server fault or an Availability Zone outage. Having a fault-tolerant architecture entails an extra cost in running additional resources than what is usually needed. This is to ensure that the mission-critical workloads are processed.

Since the scenario requires at least 2 instances to handle regular traffic, you should have 2 instances running all the time even if an AZ outage occurred. You can use an Auto Scaling Group to automatically scale your compute resources across two or more Availability Zones. You have to specify the minimum capacity to 4 instances and the maximum capacity to 6 instances. If each AZ has 2 instances running, even if an AZ fails, your system will still run a minimum of 2 instances.

Hence, the correct answer in this scenario is: Create an Auto Scaling group of EC2 instances and set the minimum capacity to 4 and the maximum capacity to 6. Deploy 2 instances in Availability Zone A and another 2 instances in

Availability Zone B.

The option that says: Create an Auto Scaling group of EC2 instances and set the minimum capacity to 2 and the maximum capacity to 6. Deploy 4 instances in Availability Zone A is incorrect because the instances are only deployed in a single Availability Zone. It cannot protect your applications and data from datacenter or AZ failures.

The option that says: Create an Auto Scaling group of EC2 instances and set the minimum capacity to 2 and the maximum capacity to 6. Use 2 Availability Zones and deploy 1 instance for each AZ is incorrect. It is required to have 2 instances running all the time. If an AZ outage happened, ASG will launch a new

instance on the unaffected AZ. This provisioning does not happen instantly, which means that for a certain period of time, there will only be 1 running instance left.

The option that says: Create an Auto Scaling group of EC2 instances and set the minimum capacity to 2 and the maximum capacity to 4. Deploy 2 instances in Availability Zone A and 2 instances in Availability Zone B is incorrect. Although this fulfills the requirement of at least 2 EC2 instances and high availability, the maximum capacity setting is wrong. It should be set to 6 to properly handle the peak load. If an AZ outage occurs and the system is at its peak load, the number of running instances in this setup will only be 4 instead of 6 and this will affect the performance of your application. References:

https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html

https://docs.aws.amazon.com/documentdb/latest/developerguide/regions-and-azs.html

Check out this AWS Auto Scaling Cheat Sheet:

https://tutorialsdojo.com/aws-auto-scaling/

Question: 86

A company is using Amazon S3 to store frequently accessed data. When an object is created or deleted, the S3 bucket will send an event notification to the Amazon SQS queue. A solutions architect needs to create a solution that will notify the development and operations team about the created or deleted objects.

Which of the following would satisfy this requirement?
A. Create an Amazon SNS topic and configure two Amazon SQS queues to subscribe to the topic. Grant Amazon S3 permission to send notifications to Amazon SNS and update the bucket to use the new SNS topic.
B. Create a new Amazon SNS FIFO topic for the other team. Grant Amazon S3 permission to send the notification to the second SNS topic.
C. Set up an Amazon SNS topic and configure two Amazon SQS queues to poll the SNS topic. Grant Amazon S3 permission to send notifications to Amazon SNS and update the bucket to use the new SNS topic.
D. Set up another Amazon SQS queue for the other team. Grant Amazon S3 permission to send a notification to the second SQS queue.

Answer: A

Explanation:

The Amazon S3 notification feature enables you to receive notifications when certain events happen in your bucket. To enable notifications, you must first add a notification configuration that identifies the events you want Amazon S3 to publish and the destinations where you want Amazon S3 to send the notifications. You store this configuration in the notification subresource that is associated with a bucket. Amazon S3 supports the following destinations where it can

publish events:

– Amazon Simple Notification Service (Amazon SNS) topic

– Amazon Simple Queue Service (Amazon SQS) queue

– AWS Lambda

In Amazon SNS, the fanout scenario is when a message published to an SNS topic is replicated and pushed to multiple endpoints, such as Amazon SQS queues, HTTP(S) endpoints, and Lambda functions. This allows for parallel asynchronous processing.
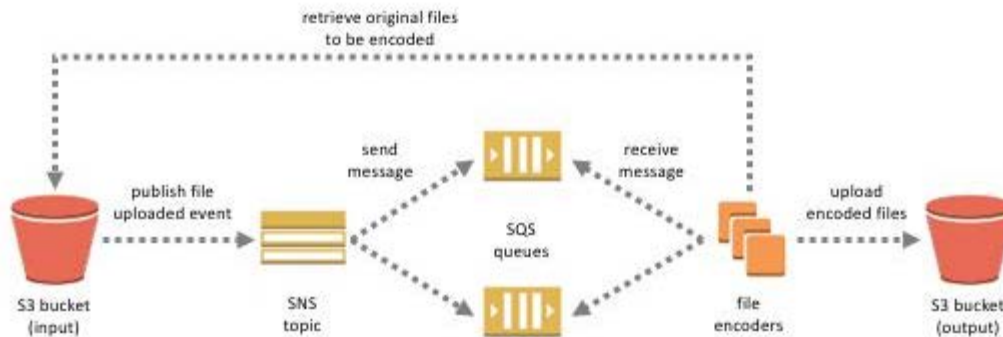
For example, you can develop an application that publishes a message to an SNS topic whenever an order is placed for a product. Then, SQS queues that are subscribed to the SNS topic receive identical notifications for the new order. An Amazon Elastic Compute Cloud (Amazon EC2) server instance attached to one of the SQS queues can handle the processing or fulfillment of the order. And you can attach another Amazon EC2 server instance to a data warehouse for analysis of all orders received. Based on the given scenario, the existing setup sends the event notification to an SQS queue. Since you need to send the notification to the development and operations team, you can use a combination of Amazon SNS and SQS. By using the message fanout pattern, you can create a topic and use two Amazon SQS queues to subscribe to the topic. If Amazon SNS receives an event notification, it will

publish the message to both subscribers.

Take note that Amazon S3 event notifications are designed to be delivered at least once and to one destination only. You cannot attach two or more SNS topics or SQS queues for S3 event notification. Therefore, you must send the event notification to Amazon SNS.

Hence, the correct answer is: Create an Amazon SNS topic and configure two Amazon SQS queues to subscribe to the topic. Grant Amazon S3 permission to send notifications to Amazon SNS and update the bucket to use the new SNS topic.

The option that says: Set up another Amazon SQS queue for the other team. Grant Amazon S3 permission to send a notification to the second SQS queue is incorrect because you can only add 1 SQS or SNS at a time for Amazon S3 events notification. If you need to send the events to multiple subscribers, you should implement a message fanout pattern with Amazon SNS and Amazon SQS.

The option that says: Create a new Amazon SNS FIFO topic for the other team. Grant Amazon S3 permission to send the notification to the second SNS topic is incorrect. Just as mentioned in the previous option, you can only add 1 SQS or SNS at a time for Amazon S3 events notification. In addition, neither Amazon SNS FIFO topic nor Amazon SQS FIFO queue is warranted in this scenario. Both of them can be used together to provide strict message ordering and

message deduplication. The FIFO capabilities of each of these services work together to act as a fully managed service to integrate distributed applications that require data consistency in near-real-time.

The option that says: Set up an Amazon SNS topic and configure two Amazon SQS queues to poll the

SNS topic. Grant Amazon S3 permission to send notifications to Amazon SNS and update the bucket to

use the new SNS topic is incorrect because you can't poll Amazon SNS. Instead of configuring queues

to poll Amazon SNS, you should configure each Amazon SQS queue to subscribe to the SNS topic.

References:

https://docs.aws.amazon.com/AmazonS3/latest/dev/ways-to-add-notification-config-to-bucket.html

https://docs.aws.amazon.com/AmazonS3/latest/dev/NotificationHowTo.html#notification-how-to-overvie

w

https://docs.aws.amazon.com/sns/latest/dg/welcome.html

Check out this Amazon S3 Cheat Sheet:

https://tutorialsdojo.com/amazon-s3/

Amazon SNS Overview:

https://youtu.be/ft5R45lEUJ8

Question: 87

An accounting application uses an RDS database configured with Multi-AZ deployments to improve availability .

What would happen to RDS if the primary database instance fails?
A. The IP address of the primary DB instance is switched to the standby DB instance.
B. The primary database instance will reboot.
C. A new database instance is created in the standby Availability Zone.
D. The canonical name record (CNAME) is switched from the primary to standby instance.

Answer: D

Explanation:

In Amazon RDS, failover is automatically handled so that you can resume database operations as quickly as possible without administrative intervention in the event that your primary database instance goes down. When failing over, Amazon RDS simply flips the canonical name record (CNAME) for your DB instance to point at the standby, which is in turn promoted to become the new primary.
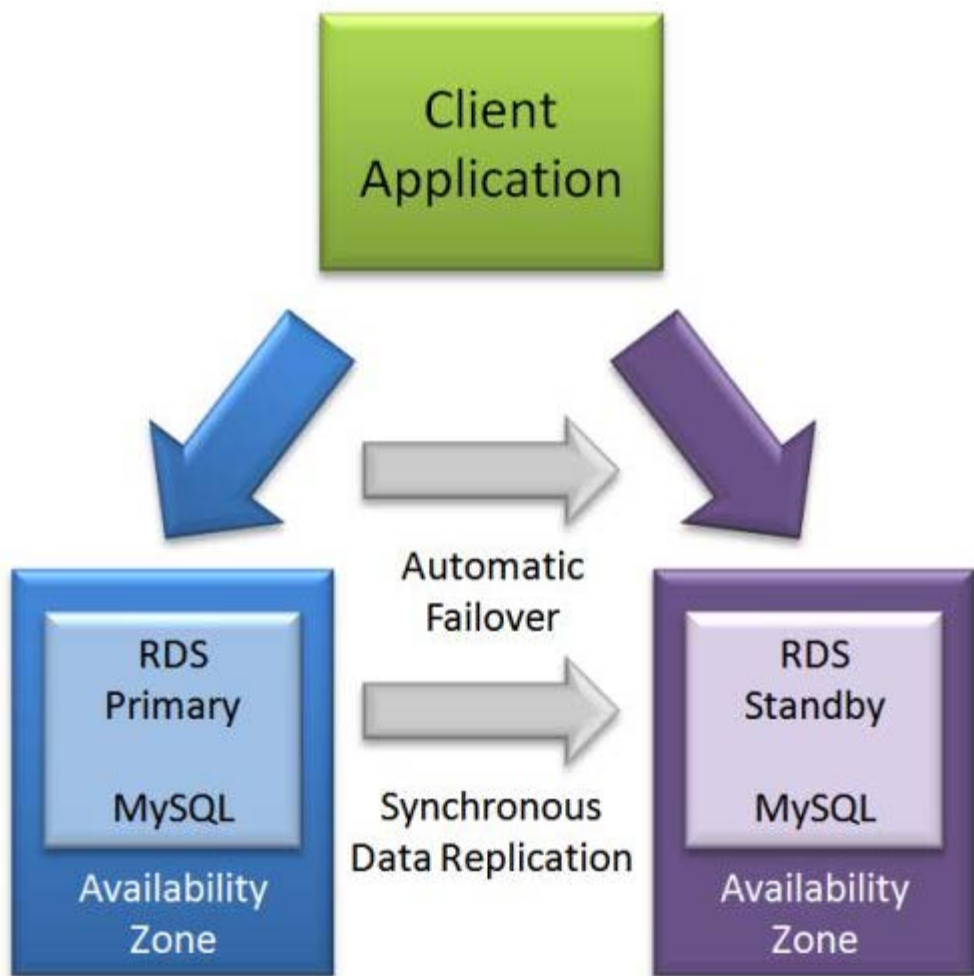
The option that says: The IP address of the primary DB instance is switched to the standby DB instance is incorrect since IP addresses are per subnet, and subnets cannot span multiple AZs.

The option that says: The primary database instance will reboot is incorrect since in the event of a failure, there is no database to reboot with.

The option that says: A new database instance is created in the standby Availability Zone is incorrect

since with multi-AZ enabled, you already have a standby database in another AZ.

References:

https://aws.amazon.com/rds/details/multi-az/

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

Amazon RDS Overview:

https://youtu.be/aZmpLl8K1UU

Check out this Amazon RDS Cheat Sheet:

https://tutorialsdojo.com/amazon-relational-database-service-amazon-rds/

Question: 88

A car dealership website hosted in Amazon EC2 stores car listings in an Amazon Aurora database managed by Amazon RDS. Once a vehicle has been sold, its data must be removed from the current listings and forwarded to a distributed processing system.

Which of the following options can satisfy the given requirement?
A. Create an RDS event subscription and send the notifications to Amazon SQ
B. Configure the SQS queues to fan out the event notifications to multiple Amazon SNS topics. Process the data using Lambda functions.
C. Create an RDS event subscription and send the notifications to AWS Lambda. Configure the Lambda function to fan out the event notifications to multiple Amazon SQS queues to update the processing system.
D. Create a native function or a stored procedure that invokes a Lambda function. Configure the Lambda function to send event notifications to an Amazon SQS queue for the processing system to consume.
E. Create an RDS event subscription and send the notifications to Amazon SN
F. Configure the SNS topic to fan out the event notifications to multiple Amazon SQS queues. Process the data using Lambda
functions.

Answer: C

Explanation:

You can invoke an AWS Lambda function from an Amazon Aurora MySQL-Compatible Edition DB cluster with a native function or a stored procedure. This approach can be useful when you want to integrate your database running on Aurora MySQL with other AWS services. For example, you might want to capture data changes whenever a row in a table is modified in your database.
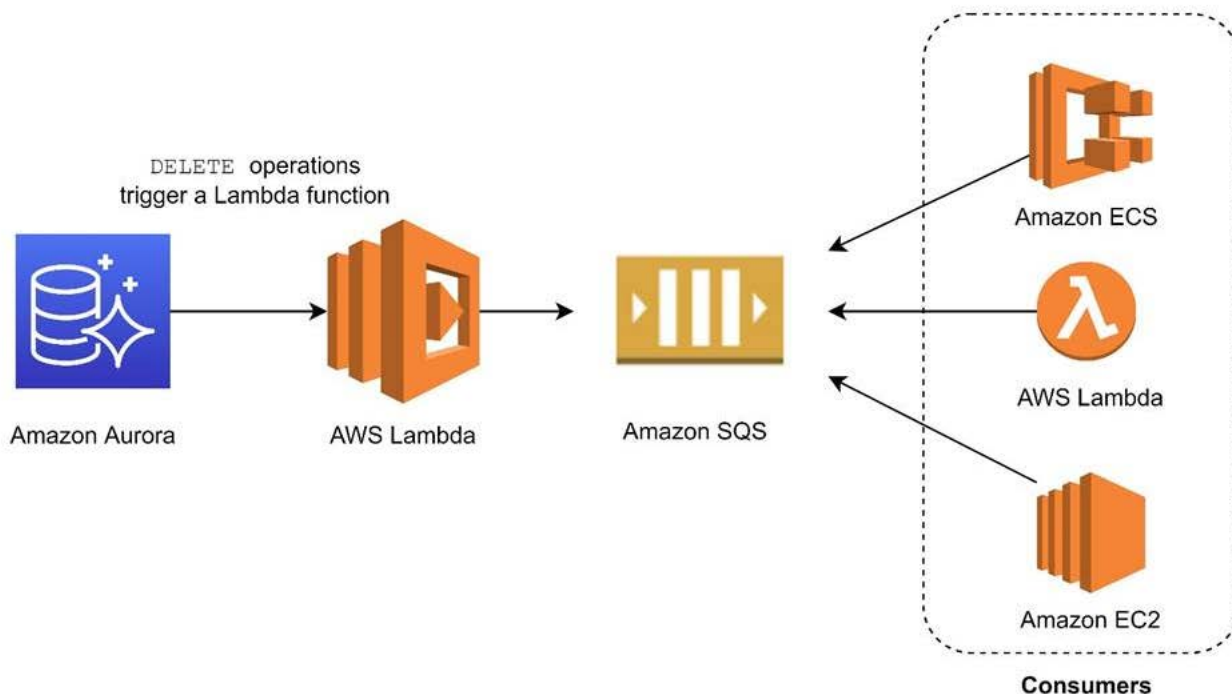
In the scenario, you can trigger a Lambda function whenever a listing is deleted from the database. You can then write the logic of the function to send the listing data to an SQS queue and have different processes consume it.

Hence, the correct answer is: Create a native function or a stored procedure that invokes a Lambda function. Configure the Lambda function to send event notifications to an Amazon SQS queue for the processing system to consume.

RDS events only provide operational events such as DB instance events, DB parameter group events, DB security group events, and DB snapshot events .

What we need in the scenario is to capture data-modifying events (INSERT, DELETE, UPDATE) which can be achieved thru native functions or stored procedures. Hence, the following options are incorrect:

– Create an RDS event subscription and send the notifications to Amazon SQS. Configure the SQS queues to fan out the event notifications to multiple Amazon SNS topics. Process the data using Lambda functions.

– Create an RDS event subscription and send the notifications to AWS Lambda. Configure the Lambda function to fan out the event notifications to multiple Amazon SQS queues to update the processing system.

– Create an RDS event subscription and send the notifications to Amazon SNS. Configure the SNS topic to fan out the event notifications to multiple Amazon SQS queues. Process the data using Lambda functions.

References:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.h tml
https://aws.amazon.com/blogs/database/capturing-data-changes-in-amazon-aurora-using-aws-lambda/ Amazon Aurora Overview:

https://youtu.be/iwS1h7rLNBQ

Check out this Amazon Aurora Cheat Sheet: https://tutorialsdojo.com/amazon-aurora/