Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

NEW QUESTION 1

An Amazon Simple Queue Service (Amazon SQS) queue serves as an event source for an AWS Lambda function In the SQS queue, each item corresponds to a video file that the Lambda function must convert to a smaller resolution The Lambda function is timing out on longer video files, but the Lambda function's timeout is already configured to its maximum value

What should a developer do to avoid the timeouts without additional code changes'?

A. Increase the memory configuration of the Lambda function
B. Increase the visibility timeout on the SQS queue
C. Increase the instance size of the host that runs the Lambda function.
D. Use multi-threading for the conversion.

Answer: A

Explanation:

Increasing the memory configuration of the Lambda function will also increase the CPU and network throughput available to the function. This can improve the performance of the video conversion process and reduce the execution time of the function. This solution does not require any code changes or additional resources. It is also recommended to follow the best practices for preventing Lambda function timeouts1. References
? Troubleshoot Lambda function invocation timeout errors | AWS re:Post

NEW QUESTION 2

An online food company provides an Amazon API Gateway HTTP API 1o receive orders for partners. The API is integrated with an AWS Lambda function. The Lambda function stores the orders in an Amazon DynamoDB table.
The company expects to onboard additional partners Some to me panthers require additional Lambda function to receive orders. The company has created an Amazon S3 bucket. The company needs 10 store all orders and updates m the S3 bucket for future analysis
How can the developer ensure that an orders and updates are stored to Amazon S3 with the LEAST development effort?

A. Create a new Lambda function and a new API Gateway API endpoin
B. Configure the new Lambda function to write to the S3 bucke
C. Modify the original Lambda function to post updates to the new API endpoint.
D. Use Amazon Kinesis Data Streams to create a new data strea
E. Modify the Lambda function to publish orders to the oats stream Configure in data stream to write to the S3 bucket.
F. Enable DynamoDB Streams on me DynamoOB tabl
G. Create a new lambda functio

Configure the Lambda function to write to the S3
H. Associate the stream's Amazon Resource Name (ARN) with the Lambda Function
bucket as records appear in the table's stream.
I. Modify the Lambda function to punish to a new Amazo
J. Simple Lambda function receives order
K. Subscribe a new Lambda function to the topi
L. Configure the new Lambda function to write to the S3 bucket as updates come through the topic.

Answer: C

Explanation:

This solution will ensure that all orders and updates are stored to Amazon S3 with the least development effort because it uses DynamoDB Streams to capture changes in the DynamoDB table and trigger a Lambda function to write those changes to the S3 bucket. This way, the original Lambda function and API Gateway API endpoint do not need to be modified, and no additional services are required. Option A is not optimal because it will require more development effort to create a new Lambda function and a new API Gateway API endpoint, and to modify the original Lambda function to post updates to the new API endpoint. Option B is not optimal because it will introduce additional costs and complexity to use Amazon Kinesis Data Streams to create a new data stream, and to modify the Lambda function to publish orders to the data stream. Option D is not optimal because it will require more development effort to modify the Lambda function to publish to a new Amazon SNS topic, and to create and subscribe a new Lambda function to the topic. References: Using DynamoDB Streams, Using AWS Lambda with Amazon S3

NEW QUESTION 3

A company notices that credentials that the company uses to connect to an external software as a service (SaaS) vendor are stored in a configuration file as plaintext.
The developer needs to secure the API credentials and enforce automatic credentials rotation on a quarterly basis.
Which solution will meet these requirements MOST securely?

A. Use AWS Key Management Service (AWS KMS) to encrypt the configuration fil
B. Decrypt the configuration file when users make API calls to the SaaS vendo
C. Enable rotation.
D. Retrieve temporary credentials from AWS Security Token Service (AWS STS) every 15 minute
E. Use the temporary credentials when users make API calls to the SaaS vendor.
F. Store the credentials in AWS Secrets Manager and enable rotatio
G. Configure the API to have Secrets Manager access.
Store the credentials in AWS Systems Manager Parameter Store and enable rotatio
H. Retrieve the credentials when users make API calls to the SaaS vendor.

Answer: C

Explanation:

Store the credentials in AWS Secrets Manager and enable rotation. Configure the API to have Secrets Manager access. This is correct. This solution will meet the requirements most securely, because it uses a service that is designed to store and manage secrets such as API credentials. AWS Secrets Manager helps you protect access to your applications, services, and IT resources by enabling you to rotate, manage, and retrieve secrets throughout their lifecycle1. You can store secrets such as passwords, database strings, API keys, and license codes as encrypted values2. You can also configure automatic rotation of your secrets on a schedule that you specify3. You can use the AWS SDK or CLI to retrieve secrets from Secrets Manager when you need them4. This way, you can avoid storing credentials in plaintext files or hardcoding them in your code.

NEW QUESTION 4
A developer is troubleshooting an Amazon API Gateway API Clients are receiving HTTP 400 response errors when the clients try to access an endpoint of the API.
How can the developer determine the cause of these errors?

A. Create an Amazon Kinesis Data Firehose delivery stream to receive API call logs from API Gatewa
B. Configure Amazon CloudWatch Logs as the delivery stream's destination.
C. Turn on AWS CloudTrail Insights and create a trail Specify the Amazon Resource Name (ARN) of the trail for the stage of the API.
D. Turn on AWS X-Ray for the API stage Create an Amazon CtoudWalch Logs log group Specify the Amazon Resource Name (ARN) of the log group for the API stage.
E. Turn on execution logging and access logging in Amazon CloudWatch Logs for the API stag
F. Create a CloudWatch Logs log grou
G. Specify the Amazon Resource Name (ARN) of the log group for the API stage.

Answer: D

Explanation:
 This solution will meet the requirements by using Amazon CloudWatch Logs to capture and analyze the logs from API Gateway. Amazon CloudWatch Logs is a service that monitors, stores, and accesses log files from AWS resources. The developer can turn on execution logging and access logging in Amazon CloudWatch Logs for the API stage, which enables logging information about API execution and client access to the API. The developer can create a CloudWatch Logs log group, which is a collection of log streams that share the same retention, monitoring, and access control settings. The developer can specify the Amazon Resource Name (ARN) of the log group for the API stage, which instructs API Gateway to send the logs to the specified log group. The developer can then examine the logs to determine the cause of the HTTP 400 response errors. Option A is not optimal because it will create an Amazon Kinesis Data Firehose delivery stream to receive API call logs from API Gateway, which may introduce additional costs and complexity for delivering and processing streaming data. Option B is not optimal because it will turn on AWS CloudTrail Insights and create a trail, which is a feature that helps identify and troubleshoot unusual API activity or operational issues, not HTTP response errors. Option C is not optimal because it will turn on AWS X-Ray for the API stage, which is a service that helps analyze and debug distributed applications, not HTTP response errors. References: [Setting Up CloudWatch Logging for a REST API], [CloudWatch Logs Concepts]

NEW QUESTION 5
A company needs to deploy all its cloud resources by using AWS CloudFormation templates A developer must create an Amazon Simple Notification Service (Amazon SNS) automatic notification to help enforce this rule. The developer creates an SNS topic and subscribes the email address of the company's security team to the SNS topic.
The security team must receive a notification immediately if an 1AM role is created without the use of CloudFormation.
Which solution will meet this requirement?

A. Create an AWS Lambda function to filter events from CloudTrail if a role was created without CloudFormation Configure the Lambda function to publish to the SNS topi
B. Create an Amazon EventBridge schedule to invoke the Lambda function every 15 minutes
C. Create an AWS Fargate task in Amazon Elastic Container Service (Amazon ECS) to filter events from CloudTrail if a role was created without CloudFormation Configure the Fargate task to publish to the SNS topic Create an Amazon EventBridge schedule to run the Fargate task every 15 minutes
D. Launch an Amazon EC2 instance that includes a script to filter events from CloudTrail if a role was created without CloudFormatio
E. Configure the script to publish to the SNS topi
F. Create a cron job to run the script on the EC2 instance every 15 minutes.
G. Create an Amazon EventBridge rule to filter events from CloudTrail if a role was created without CloudFormation Specify the SNS topic as the target of the EventBridge rule.

Answer: D

Explanation:
 Creating an Amazon EventBridge rule is the most efficient and scalable way to monitor and react to events from CloudTrail, such as the creation of an IAM role without CloudFormation. EventBridge allows you to specify a filter pattern to match the events you are interested in, and then specify an SNS topic as the target to send notifications. This solution does not require any additional resources or code, and it can trigger notifications in near real-time. The other solutions involve creating and managing additional resources, such as Lambda functions, Fargate tasks, or EC2 instances, and they rely on polling CloudTrail events every 15 minutes, which can introduce delays and increase
costs. References
? Using Amazon EventBridge rules to process AWS CloudTrail events
? Using AWS CloudFormation to create and manage AWS Batch resources
? How to use AWS CloudFormation to configure auto scaling for Amazon Cognito and AWS AppSync
? Using AWS CloudFormation to automate the creation of AWS WAF web ACLs, rules, and conditions

NEW QUESTION 6
An application that runs on AWS receives messages from an Amazon Simple Queue Service (Amazon SQS) queue and processes the messages in batches. The application sends the data to another SQS queue to be consumed by another legacy application. The legacy system can take up to 5 minutes to process some transaction data.
A developer wants to ensure that there are no out-of-order updates in the legacy system. The developer cannot alter the behavior of the legacy system.
Which solution will meet these requirements?

A. Use an SQS FIFO queue
B. Configure the visibility timeout value.
C. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ
D. Configure the DelaySeconds values.
E. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ
F. Configure the visibility timeout value.
G. Use an SQS FIFO queu
H. Configure the DelaySeconds value.

Answer: A

Explanation:
? An SQS FIFO queue is a type of queue that preserves the order of messages and ensures that each message is delivered and processed only once1. This is suitable for the scenario where the developer wants to ensure that there are no out-of-order updates in the legacy system.
? The visibility timeout value is the amount of time that a message is invisible in the queue after a consumer receives it2. This prevents other consumers from

processing the same message simultaneously. If the consumer does not delete the message before the visibility timeout expires, the message becomes visible again and another consumer can receive it2.
? In this scenario, the developer needs to configure the visibility timeout value to be longer than the maximum processing time of the legacy system, which is 5 minutes. This will ensure that the message remains invisible in the queue until the legacy system finishes processing it and deletes it. This will prevent duplicate or out-of-order processing of messages by the legacy system.

NEW QUESTION 7
A company has an application that is hosted on Amazon EC2 instances The application stores objects in an Amazon S3 bucket and allows users to download objects from the S3 bucket A developer turns on S3 Block Public Access for the S3 bucket After this change, users report errors when they attempt to download objects The developer needs to implement a solution so that only users who are signed in to the application can access objects in the S3 bucket.
Which combination of steps will meet these requirements in the MOST secure way? (Select TWO.)

A. Create an EC2 instance profile and role with an appropriate policy Associate the role with the EC2 instances
B. Create an 1AM user with an appropriate polic
C. Store the access key ID and secret access key on the EC2 instances
D. Modify the application to use the S3 GeneratePresignedUrl API call
E. Modify the application to use the S3 GetObject API call and to return the object handle to the user
F. Modify the application to delegate requests to the S3 bucket.

Answer: AC

Explanation:
The most secure way to allow the EC2 instances to access the S3 bucket is to use an EC2 instance profile and role with an appropriate policy that grants the necessary permissions. This way, the EC2 instances can use temporary security credentials that are automatically rotated and do not need to store any access keys on the instances. To allow the users who are signed in to the application to download objects from the S3 bucket, the application can use the S3 GeneratePresignedUrl API call to create a pre-signed URL that grants temporary access to a specific object. The pre-signed URL can be returned to the user, who can then use it to download the object within a specified time period. References
? Use Amazon S3 with Amazon EC2
? How to Access AWS S3 Bucket from EC2 Instance In a Secured Way
? Sharing an Object with Others

NEW QUESTION 8
A company runs a payment application on Amazon EC2 instances behind an Application Load Balance The EC2 instances run in an Auto Scaling group across multiple Availability Zones The application needs to retrieve application secrets during the application startup and export the secrets as environment variables These secrets must be encrypted at rest and need to be rotated every month.
Which solution will meet these requirements with the LEAST development effort?

A. Save the secrets in a text file and store the text file in Amazon S3 Provision a customer managed key Use the key for secret encryption in Amazon S3 Read the contents of the text file and read the export as environment variables Configure S3 Object Lambda to rotate the text file every month
B. Save the secrets as strings in AWS Systems Manager Parameter Store and use the default AWS Key Management Service (AWS KMS) key Configure an Amazon EC2 user data script to retrieve the secrets during the startup and export as environment variables Configure an AWS Lambda function to rotate the secrets in Parameter Store every month.
C. Save the secrets as base64 encoded environment variables in the application propertie
D. Retrieve the secrets during the application startu
E. Reference the secrets in the application cod
F. Write a script to rotate the secrets saved as environment variables.
G. Store the secrets in AWS Secrets Manager Provision a new customer master key Use the key to encrypt the secrets Enable automatic rotation Configure an Amazon EC2 user data script to programmatically retrieve the secrets during the startup and export as environment variables

Answer: D

Explanation:
AWS Secrets Manager is a service that enables the secure management and rotation of secrets, such as database credentials, API keys, or passwords. By using Secrets Manager, the company can avoid hardcoding secrets in the application code or properties files, and instead retrieve them programmatically during the application startup. Secrets Manager also supports automatic rotation of secrets by using AWS Lambda functions or built-in rotation templates. The company can provision a customer master key (CMK) to encrypt the secrets and use the AWS SDK or CLI to export the secrets as environment variables. References:
? What Is AWS Secrets Manager? - AWS Secrets Manager
? Rotating Your AWS Secrets Manager Secrets - AWS Secrets Manager
? Retrieving a Secret - AWS Secrets Manager

NEW QUESTION 9
A company uses Amazon API Gateway to expose a set of APIs to customers. The APIs have caching enabled in API Gateway. Customers need a way to invalidate the cache for each API when they test the API.
What should a developer do to give customers the ability to invalidate the API cache?

A. Ask the customers to use AWS credentials to call the InvalidateCache API operation.
B. Attach an InvalidateCache policy to the IAM execution role that the customers use to invoke the AP
C. Ask the customers to send a request that contains the HTTP header when they make an API call.
D. Ask the customers to use the AWS SDK API Gateway class to invoke the InvalidateCache API operation.
E. Attach an InvalidateCache policy to the IAM execution role that the customers use to invoke the AP
F. Ask the customers to add the INVALIDATE_CACHE query string parameter when they make an API call.

Answer: D

NEW QUESTION 10
A developer is creating a simple proof-of-concept demo by using AWS CloudFormation and AWS Lambda functions The demo will use a CloudFormation template to deploy an existing Lambda function The Lambda function uses deployment packages and dependencies stored in Amazon S3 The developer defined anAWS Lambda Function resource in a CloudFormation template. The developer needs to add the S3 bucket to the CloudFormation template.

What should the developer do to meet these requirements with the LEAST development effort?

A. Add the function code in the CloudFormation template inline as the code property
B. Add the function code in the CloudFormation template as the ZipFile property.
C. Find the S3 key for the Lambda function Add the S3 key as the ZipFile property in the CloudFormation template.
D. Add the relevant key and bucket to the S3Bucket and S3Key properties in the CloudFormation template

Answer: D

Explanation:
The easiest way to add the S3 bucket to the CloudFormation template is to use the S3Bucket and S3Key properties of the AWS::Lambda::Function resource. These properties specify the name of the S3 bucket and the location of the .zip file that contains the function code and dependencies. This way, the developer function code or upload it to a different location. The other options are either not feasible or not efficient. does not need to modify the

The code property can only be used for inline code, not for code stored in S3. The ZipFile property can only be used for code that is less than 4096 bytes, not for code that has dependencies. Finding the S3 key for the Lambda function and adding it as the ZipFile property would not work, as the ZipFile property expects a base64-encoded .zip file, not an S3 location. References
? AWS::Lambda::Function - AWS CloudFormation
? Deploying Lambda functions as .zip file archives
? AWS Lambda Function Code - AWS CloudFormation

NEW QUESTION 10
A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream. The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is increasing and that the Lambda run duration is constantly above normal.
Which actions should the developer take to increase the processing speed? (Choose two.)

A. Increase the number of shards of the Kinesis data stream.
B. Decrease the timeout of the Lambda function.
C. Increase the memory that is allocated to the Lambda function.
D. Decrease the number of shards of the Kinesis data stream.
E. Increase the timeout of the Lambda function.

Answer: AC

Explanation:
Increasing the number of shards of the Kinesis data stream will increase the throughput and parallelism of the data processing. Increasing the memory that is allocated to the Lambda function will also increase the CPU and network performance of the function, which will reduce the run duration and improve the processing speed. Option B is not correct because decreasing the timeout of the Lambda function will not affect the processing speed, but may cause some records to fail if they exceed the timeout limit. Option D is not correct because decreasing the number of shards of the Kinesis data stream will decrease the throughput and parallelism of the data processing, which will slow down the processing speed. Option E is not correct because increasing the timeout of the Lambda function will not affect the processing speed, but may increase the cost of running the function.
References: [Amazon Kinesis Data Streams Scaling], [AWS Lambda Performance Tuning]

NEW QUESTION 13
A developer is creating an AWS CloudFormation template to deploy Amazon EC2 instances across multiple AWS accounts. The developer must choose the EC2 instances from a list of approved instance types.
How can the developer incorporate the list of approved instance types in the CloudFormation template?

A. Create a separate CloudFormation template for each EC2 instance type in the list.
B. In the Resources section of the CloudFormation template, create resources for each EC2 instance type in the list.
C. In the CloudFormation template, create a separate parameter for each EC2 instance type in the list.
D. In the CloudFormation template, create a parameter with the list of EC2 instance types as AllowedValues.

Answer: D

Explanation:
In the CloudFormation template, the developer should create a parameter with the list of approved EC2 instance types as AllowedValues. This way, users can select the instance type they want to use when launching the CloudFormation stack, but only from the approved list.

NEW QUESTION 18
A developer is working on a Python application that runs on Amazon EC2 instances. The developer wants to enable tracing of application requests to debug performance issues in the code.
Which combination of actions should the developer take to achieve this goal? (Select TWO)

A. Install the Amazon CloudWatch agent on the EC2 instances.
B. Install the AWS X-Ray daemon on the EC2 instances.
C. Configure the application to write JSON-formatted togs to /var/log/cloudwatch.
D. Configure the application to write trace data to /Var/log-/xray.
E. Install and configure the AWS X-Ray SDK for Python in the application.

Answer: BE

Explanation:
This solution will meet the requirements by using AWS X-Ray to enable tracing of application requests to debug performance issues in the code. AWS X-Ray is a service that collects data about requests that the applications serve, and provides tools to view, filter, and gain insights into that data.
The developer can install the AWS X-Ray daemon on the EC2 instances, which is a software that listens for traffic on UDP port 2000, gathers raw segment data, and relays it to the X-Ray API. The developer can also install and configure the AWS X-Ray SDK for Python in the application, which is a library that enables instrumenting Python code to generate and send trace data to the X-Ray daemon. Option A is not optimal because it will install the Amazon CloudWatch agent on the EC2 instances, which is a software that collects metrics and logs from EC2 instances and on- premises servers, not application performance data. Option C is not optimal because it will configure the application to write JSON-formatted logs to /var/log/cloudwatch, which is not a valid path or destination for CloudWatch logs. Option D is not optimal because it will configure the application to write trace data to /var/log/xray, which is also not a valid path or destination for X-Ray trace

data.
References: [AWS X-Ray], [Running the X-Ray Daemon on Amazon EC2]

NEW QUESTION 22
A company uses a custom root certificate authority certificate chain (Root CA Cert) that is 10 KB in size generate SSL certificates for its on-premises HTTPS endpoints. One of the company's cloud based applications has hundreds of AWS Lambda functions that pull date from these endpoints. A developer updated the trust store of the Lambda execution environment to use the Root CA Cert when the Lambda execution environment is initialized. The developer bundled the Root CA Cert as a text file in the Lambdas deployment bundle.
After 3 months of development the root CA Cert is no longer valid and must be updated. The developer needs a more efficient solution to update the Root CA Cert for all deployed Lambda functions. The solution must not include rebuilding or updating all Lambda functions that use the Root CA Cert. The solution must also work for all development, testing and production environment. Each environment is managed in a separate AWS account.
When combination of steps Would the developer take to meet these environments MOST cost-effectively? (Select TWO)

A. Mastered
B. Not Mastered

Answer: A

Explanation:
This solution will meet the requirements by storing the Root CA Cert as a Secure String parameter in AWS Systems Manager Parameter Store, which is a secure and scalable service for storing and managing configuration data and secrets. The resource-based policy will allow IAM users in different AWS accounts and environments to access the parameter without requiring cross-account roles or permissions. The Lambda code will be refactored to load the Root CA Cert from the parameter store and modify the runtime trust store outside the Lambda function handler, which will improve performance and reduce latency by avoiding repeated calls to Parameter Store and trust store modifications for each invocation of the Lambda function. Option A is not optimal because it will use AWS Secrets Manager instead of AWS Systems Manager Parameter Store, which will incur additional costs and complexity for storing and managing a non-secret configuration data such as Root CA Cert. Option C is not optimal because it will deactivate the application secrets and monitor the application error logs temporarily, which will cause application downtime and potential data loss. Option D is not optimal because it will modify the runtime trust store inside the Lambda function handler, which will degrade performance and increase latency by repeating unnecessary operations for each invocation of the Lambda function.
References: AWS Systems Manager Parameter Store, [Using SSL/TLS to Encrypt a Connection to a DB Instance]

NEW QUESTION 26
A company receives food orders from multiple partners. The company has a microservices application that uses Amazon API Gateway APIs with AWS Lambda integration. Each partner sends orders by calling a customized API that is exposed through API Gateway. The API call invokes a shared Lambda function to process the orders.
Partners need to be notified after the Lambda function processes the orders. Each partner must receive updates for only the partner's own orders. The company wants to add new partners in the future with the fewest code changes possible.
Which solution will meet these requirements in the MOST scalable way?

A. Create a different Amazon Simple Notification Service (Amazon SNS) topic for each partne
B. Configure the Lambda function to publish messages for each partner to the partner's SNS topic.
C. Create a different Lambda function for each partne
D. Configure the Lambda function to notify each partner's service endpoint directly.
E. Create an Amazon Simple Notification Service (Amazon SNS) topi
F. Configure the Lambda function to publish messages with specific attributes to the SNS topi
G. Subscribe each partner to the SNS topi
H. Apply the appropriate filter policy to the topic subscriptions.
I. Create one Amazon Simple Notification Service (Amazon SNS) topi
J. Subscribe all partners to the SNS topic.

Answer: C

Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fully managed messaging service that enables pub/sub communication between distributed systems. The developer can create an SNS topic and configure the Lambda function to publish messages with specific attributes to the topic. The developer can subscribe each partner to the SNS topic and apply the appropriate filter policy to the topic subscriptions. This way, each partner will receive updates for only their own orders based on the message attributes. This solution will meet the requirements in the most scalable way and allow adding new partners in the future with minimal code changes.
References:
? [Amazon Simple Notification Service (SNS)]
? [Filtering Messages with Attributes - Amazon Simple Notification Service]

NEW QUESTION 31
A company has an application that runs as a series of AWS Lambda functions. Each Lambda function receives data from an Amazon Simple Notification Service (Amazon SNS) topic and writes the data to an Amazon Aurora DB instance.
To comply with an information security policy, the company must ensure that the Lambda functions all use a single securely encrypted database connection string to access Aurora.
Which solution will meet these requirements'?

A. Use IAM database authentication for Aurora to enable secure database connections for ail the Lambda functions.
B. Store the credentials and read the credentials from an encrypted Amazon RDS DB instance.
C. Store the credentials in AWS Systems Manager Parameter Store as a secure string parameter.
D. Use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption.

Answer: A

Explanation:
This solution will meet the requirements by using IAM database authentication for Aurora which enables using IAM roles or users to authenticate with Aurora databases instead of using passwords or other secrets. The developer can use IAM database authentication for Aurora to enable secure database connections for all the Lambda functions that access Aurora DB instance. The developer can create an IAM role with permission to connect to Aurora DB instance and attach it to each Lambda function. The developer can also configure Aurora DB instance to use IAM database authentication and enable encryption in transit

using SSL certificates. This way, the Lambda functions can use a single securely encrypted database connection string to access Aurora without needing any secrets or passwords. Option B is not optimal because it will store the credentials and read them from an encrypted Amazon RDS DB instance, which may introduce additional costs and complexity for managing and accessing another RDS DB instance. Option C is not optimal because it will store the credentials in AWS Systems Manager Parameter Store as a secure string parameter, which may require additional steps or permissions to retrieve and decrypt the credentials from Parameter Store. Option D is not optimal because it will use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption, which may not be secure or scalable as environment variables are stored as plain text unless encrypted with AWS KMS. References: [IAM Database Authentication for MySQL and PostgreSQL], [Using SSL/TLS to Encrypt a Connection to a DB Instance]

NEW QUESTION 32
A developer must use multi-factor authentication (MFA) to access data in an Amazon S3
bucket that is in another AWS account. Which AWS Security Token Service (AWS STS) API operation should the developer use with the MFA information to meet this requirement?

A. AssumeRoleWithWebidentity
B. GetFederationToken
C. AssumeRoleWithSAML
D. AssumeRole

Answer: D

Explanation:
The AssumeRole API operation returns a set of temporary security credentials that can be used to access resources in another AWS account. The developer can specify the MFA device serial number and the MFA token code in the request parameters. This option enables the developer to use MFA to access data in an S3 bucket that is in another AWS account. The other options are not relevant or effective for this scenario. References
? AssumeRole
? Requesting Temporary Security Credentials

NEW QUESTION 37
A company is building a web application on AWS. When a customer sends a request, the application will generate reports and then make the reports available to the customer within one hour. Reports should be accessible to the customer for 8 hours. Some reports are larger than 1 MB. Each report is unique to the customer. The application should delete all reports that are older than 2 days.
Which solution will meet these requirements with the LEAST operational overhead?

A. Generate the reports and then store the reports as Amazon DynamoDB items that have a specified TT
B. Generate a URL that retrieves the reports from DynamoD
C. Provide the URL to customers through the web application.
D. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryptio
E. Attach the reports to an Amazon Simple Notification Service (Amazon SNS) messag
F. Subscribe the customer to email notifications from Amazon SNS.
G. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryptio
H. Generate a presigned URL that contains an expiration date Provide the URL to customers through the web applicatio
I. Add S3 Lifecycle configuration rules to the S3 bucket to delete old reports.
J. Generate the reports and then store the reports in an Amazon RDS database with a date stam
K. Generate an URL that retrieves the reports from the RDS databas
L. Provide the URL to customers through the web applicatio
M. Schedule an hourly AWS Lambda function to delete database records that have expired date stamps.

Answer: C

Explanation:
This solution will meet the requirements with the least operational overhead because it uses Amazon S3 as a scalable, secure, and durable storage service for the reports. The presigned URL will allow customers to access their reports for a limited time (8 hours) without requiring additional authentication. The S3 Lifecycle configuration rules will automatically delete the reports that are older than 2 days, reducing storage costs and complying with the data retention policy. Option A is not optimal because it will incur additional costs and complexity to store the reports as DynamoDB items, which have a size limit of 400 KB. Option B is not optimal because it will not provide customers with access to their reports within one hour, as Amazon SNS email delivery is not guaranteed. Option D is not optimal because it will require more operational overhead to manage an RDS database and a Lambda function for storing and deleting the reports.
References: Amazon S3 Presigned URLs, Amazon S3 Lifecycle

NEW QUESTION 42
For a deployment using AWS Code Deploy, what is the run order of the hooks for in-place deployments?

A. BeforeInstall -> ApplicationStop -> ApplicationStart -> AfterInstall
B. ApplicationStop -> BeforeInstall -> AfterInstall -> ApplicationStart
C. BeforeInstall -> ApplicationStop -> ValidateService -> ApplicationStart
D. ApplicationStop -> BeforeInstall -> ValidateService -> ApplicationStart

Answer: B

Explanation:
For in-place deployments, AWS CodeDeploy uses a set of predefined hooks that run in a specific order during each deployment lifecycle event. The hooks are ApplicationStop, BeforeInstall, AfterInstall, ApplicationStart, and ValidateService. The run order of the hooks for in-place deployments is as follows:
? ApplicationStop: This hook runs first on all instances and stops the current
application that is running on the instances.
? BeforeInstall: This hook runs after ApplicationStop on all instances and performs any tasks required before installing the new application revision.
? AfterInstall: This hook runs after BeforeInstall on all instances and performs any
tasks required after installing the new application revision.
? ApplicationStart: This hook runs after AfterInstall on all instances and starts the new application that has been installed on the instances.
? ValidateService: This hook runs last on all instances and verifies that the new application is running properly on the instances.

Reference: [AWS CodeDeploy lifecycle event hooks reference]

NEW QUESTION 43
A company runs an application on AWS The application stores data in an Amazon DynamoDB table Some queries are taking a long time to run These slow queries involve an attribute that is not the table's partition key or sort key
The amount of data that the application stores in the DynamoDB table is expected to increase significantly. A developer must increase the performance of the queries.
Which solution will meet these requirements'?

A. Increase the page size for each request by setting the Limit parameter to be higher than the default value Configure the application to retry any request that exceeds the provisioned throughput.
B. Create a global secondary index (GSI). Set query attribute to be the partition key of the index
C. Perform a parallel scan operation by issuing individual scan requests in the parameters specify the segment for the scan requests and the total number of segments for the parallel scan.
D. Turn on read capacity auto scaling for the DynamoDB tabl
E. Increase the maximum read capacity units (RCUs).

Answer: B

Explanation:
Creating a global secondary index (GSI) is the best solution to improve the performance of the queries that involve an attribute that is not the table's partition key or sort key. A GSI allows you to define an alternate key for your table and query the data using that key. This way, you can avoid scanning the entire table and reduce the latency and cost of your queries. You should also follow the best practices for designing and using GSIs in DynamoDB12. References
? Working with Global Secondary Indexes - Amazon DynamoDB
? DynamoDB Performance & Latency - Everything You Need To Know

NEW QUESTION 48
A company wants to share information with a third party. The third party has an HTTP API endpoint that the company can use to share the information. The company has the required API key to access the HTTP API.
The company needs a way to manage the API key by using code. The integration of the API key with the application code cannot affect application performance. Which solution will meet these requirements MOST securely?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
AWS Secrets Manager is a service that helps securely store, rotate, and manage secrets such as API keys, passwords, and tokens. The developer can store the API credentials in AWS Secrets Manager and retrieve them at runtime by using the AWS SDK. This solution will meet the requirements of security, code management, and performance. Storing the API credentials in a local code variable or an S3 object is not secure, as it exposes the credentials to unauthorized access or leakage. Storing the API credentials in a DynamoDB table is also not secure, as it requires additional encryption and access control measures. Moreover, retrieving the credentials from S3 or DynamoDB may affect application performance due to network latency.
References:
? [What Is AWS Secrets Manager? - AWS Secrets Manager]
? [Retrieving a Secret - AWS Secrets Manager]

NEW QUESTION 49
A developer is testing an application that invokes an AWS Lambda function asynchronously. During the testing phase the Lambda function fails to process after two retries.
How can the developer troubleshoot the failure?

A. Configure AWS CloudTrail logging to investigate the invocation failures.
B. Configure Dead Letter Queues by sending events to Amazon SQS for investigation.
C. Configure Amazon Simple Workflow Service to process any direct unprocessed events.
D. Configure AWS Config to process any direct unprocessed events.

Answer: B

Explanation:
This solution allows the developer to troubleshoot the failure by capturing unprocessed events in a queue for further analysis. Dead Letter Queues (DLQs) are queues that store messages that could not be processed by a service, such as Lambda, for various reasons, such as configuration errors, throttling limits, or permissions issues. The developer can configure DLQs for Lambda functions by sending events to either an Amazon Simple Queue Service (SQS) queue or an Amazon Simple Notification Service (SNS) topic. The developer can then inspect the messages in the queue or topic to identify and fix the root cause of the failure. Configuring AWS CloudTrail logging will not capture invocation failures for asynchronous Lambda invocations, but only record API calls made by or on behalf of Lambda. Configuring Amazon Simple Workflow Service (SWF) or AWS Config will not process any direct unprocessed events, but require additional integration and configuration.
Reference: [Using AWS Lambda with DLQs], [Asynchronous invocation]

NEW QUESTION 50
A developer is creating a new REST API by using Amazon API Gateway and AWS Lambda. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.
The developer wants to make the REST API available for testing by using API Gateway locally.
Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

A. Sam local invoke
B. Sam local generate-event
C. Sam local start-lambda
D. Sam local start-api

Answer: D

Explanation:
 The AWS Serverless Application Model Command Line Interface (AWS SAM CLI) is a command-line tool for local development and testing of Serverless applications2. The sam local start-api subcommand of AWS SAM CLI is used to simulate a REST API by starting a new local endpoint3. Therefore, option D is correct.

NEW QUESTION 55
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions. When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

A. Set the Deployment Preference Type to Canaryl OPercent10Minute
B. Set the AutoPublishAlias property to the Lambda alias.
C. Set the Deployment Preference Type to Linearl OPercentEveryIOMinute
D. Set AutoPublishAlias property to the Lambda alias.
E. Set the Deployment Preference Type to Canaryl OPercentIOMinute
F. Set the PreTraffic and PostTraffic properties to the Lambda alias.
G. Set the Deployment Preference Type to Linearl OPercentEvery10Minute
H. Set PreTraffic and PostTraffic properties to the Lambda alias.

Answer: A

Explanation:
? The Deployment Preference Type property specifies how traffic should be shifted between versions of a Lambda function1. The Canary10Percent10Minutes option means that 10% of the traffic is immediately shifted to the new version, and after 10 minutes, the remaining 90% of the traffic is shifted1. This matches the requirement of shifting 10% of the traffic for the first 10 minutes, and then switching all traffic to the new version.
? The AutoPublishAlias property enables AWS SAM to automatically create and update a Lambda alias that points to the latest version of the function1. This is required to use the Deployment Preference Type property1. The alias name can be specified by the developer, and it can be used to invoke the function with the latest code.

NEW QUESTION 60
A company wants to automate part of its deployment process. A developer needs to automate the process of checking for and deleting unused resources that supported previously deployed stacks but that are no longer used.
The company has a central application that uses the AWS Cloud Development Kit (AWS CDK) to manage all deployment stacks. The stacks are spread out across multiple accounts. The developer's solution must integrate as seamlessly as possible within the current deployment process.
Which solution will meet these requirements with the LEAST amount of configuration?

A.   In the central AWS CDK application, write a handler function in the code that uses AWS SDK calls to check for and delete unused resource
B. Create an AWS CloudPormation template from a JSON fil
C. Use the template to attach the function code to an AWS Lambda function and lo invoke the Lambda function when the deployment slack runs.
D. In the central AWS CDK applicatio
E. write a handler function in the code that uses AWS SDK calls to check for and delete unused resource
F. Create an AWS CDK custom resource Use the custom resource to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.
G. In the central AWS CDK, write a handler function m the code that uses AWS SDK calls to check for and delete unused resource
H. Create an API in AWS Amplify Use the API to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.
I. In the AWS Lambda console write a handler function in the code that uses AWS SDK calls to check for and delete unused resource
J. Create an AWS CDK custom resourc
K. Use the custom resource to import the Lambda function into the stack and to Invoke the Lambda function when the deployment stack runs.

Answer: B

Explanation:
 This solution meets the requirements with the least amount of configuration because it uses a feature of AWS CDK that allows custom logic to be executed during stack deployment or deletion. The AWS Cloud Development Kit (AWS CDK) is a software development framework that allows you to define cloud infrastructure as code and provision it through CloudFormation. An AWS CDK custom resource is a construct that enables you to create resources that are not natively supported by CloudFormation or perform tasks that are not supported by CloudFormation during stack deployment or deletion. The developer can write a handler function in the code that uses AWS SDK calls to check for and delete unused resources, and create an AWS CDK custom resource that attaches the function code to a Lambda function and invokes it when the deployment stack runs. This way, the developer can automate the cleanup process without requiring additional configuration or integration. Creating a CloudFormation template from a JSON file will require additional configuration and integration with the central AWS CDK application. Creating an API in AWS Amplify will require additional configuration and integration with the central AWS CDK application and may not provide optimal performance or availability. Writing a handler function in the AWS Lambda console will require additional configuration and integration with the central AWS CDK application.
Reference: [AWS Cloud Development Kit (CDK)], [Custom Resources]

NEW QUESTION 65
An application uses Lambda functions to extract metadata from files uploaded to an S3 bucket; the metadata is stored in Amazon DynamoDB. The application starts behaving unexpectedly, and the developer wants to examine the logs of the Lambda function code for errors.
Based on this system configuration, where would the developer find the logs?

A. Amazon S3
B. AWS CloudTrail
C. Amazon CloudWatch
D. Amazon DynamoDB

Answer: C

Explanation:
Amazon CloudWatch is the service that collects and stores logs from AWS Lambda functions. The developer can use CloudWatch Logs Insights to query and analyze the logs for errors and metrics. Option A is not correct because Amazon S3 is a storage service that does not store Lambda function logs. Option B is not correct because AWS CloudTrail is a service that records API calls and events for AWS services, not Lambda function logs. Option D is not correct because Amazon DynamoDB is a database service that does not store Lambda function logs.
References: AWS Lambda Monitoring, [CloudWatch Logs Insights]

NEW QUESTION 68
A developer at a company needs to create a small application mat makes the same API call once each flay at a designated time. The company does not have infrastructure in the AWS Cloud yet, but the company wants to implement this functionality on AWS.
Which solution meets these requirements in the MOST operationally efficient manner?

A. Use a Kubermetes cron job that runs on Amazon Elastic Kubemetes Sen/ice (Amazon EKS)
B. Use an Amazon Linux crontab scheduled job that runs on Amazon EC2
C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event.
D. Use an AWS Batch job that is submitted to an AWS Batch job queue.

Answer: C

Explanation:
This solution meets the requirements in the most operationally efficient manner because it does not require any infrastructure provisioning or management. The developer can create a Lambda function that makes the API call and configure an EventBridge rule that triggers the function once a day at a designated time. This is a serverless solution that scales automatically and only charges for the execution time of the function.
Reference: [Using AWS Lambda with Amazon EventBridge], [Schedule Expressions for
                        Rules]

NEW QUESTION 73
A developer has written an AWS Lambda function. The function is CPU-bound. The developer wants to ensure that the function returns responses quickly.
How can the developer improve the function's performance?

A. Increase the function's CPU core count.
B. Increase the function's memory.
C. Increase the function's reserved concurrency.
D. Increase the function's timeout.

Answer: B

Explanation:
The amount of memory you allocate to your Lambda function also determines how much CPU and network bandwidth it gets. Increasing the memory size can improve the performance of CPU-bound functions by giving them more CPU power. The CPU allocation is proportional to the memory allocation, so a function with 1 GB of memory has twice the CPU power of a function with 512 MB of memory. Reference: AWS Lambda execution environment

NEW QUESTION 75
A company hosts a client-side web application for one of its subsidiaries on Amazon S3. The web application can be accessed through Amazon CloudFront from https://www.example.com. After a successful rollout, the company wants to host three more client-side web applications for its remaining subsidiaries on three separate S3 buckets.
To achieve this goal, a developer moves all the common JavaScript files and web fonts to a central S3 bucket that serves the web applications. However, during testing, the developer notices that the browser blocks the JavaScript files and web fonts.
What should the developer do to prevent the browser from blocking the JavaScript files and web fonts?

A. Create four access points that allow access to the central S3 bucke
B. Assign an access point to each web application bucket.
C. Create a bucket policy that allows access to the central S3 bucke
D. Attach the bucket policy to the central S3 bucket.
E. Create a cross-origin resource sharing (CORS) configuration that allows access to the central S3 bucke
F. Add the CORS configuration to the central S3 bucket.
G. Create a Content-MD5 header that provides a message integrity check for the central S3 bucke
H. Insert the Content-MD5 header for each web application request.

Answer: C

Explanation:
This is a frequent trouble. Web applications cannot access the resources in other domains by default, except some exceptions. You must configure CORS on the resources to be accessed. https://docs.aws.amazon.com/AmazonS3/latest/userguide/cors.html

NEW QUESTION 77
A company has multiple Amazon VPC endpoints in the same VPC. A developer needs configure an Amazon S3 bucket policy so users can access an S3 bucket only by using these VPC endpoints.
Which solution will meet these requirements?

A. Create multiple S3 bucket polices by using each VPC endpoint ID that have the aws SourceVpce value in the StringNotEquals condition.
B. Create a single S3 bucket policy that has the aws SourceVpc value and in the StingNotEquals condition to use VPC ID.
C. Create a single S3 bucket policy that the multiple aws SourceVpce value and in the SringNotEquals condton to use vpce.
D. Create a single S3 bucket policy that has multiple aws sourceVpce value in the StingNotEquale conditio
E. Repeat for all the VPC endpoint IDs.

Answer: D

Explanation:
This solution will meet the requirements by creating a single S3 bucket policy that denies access to the S3 bucket unless the request comes from one of the specified VPC endpoints. The aws:SourceVpce condition key is used to match the ID of the VPC endpoint that is used to access the S3 bucket. The

StringNotEquals condition operator is used to negate the condition, so that only requests from the listed VPC endpoints are allowed.

Option A is not optimal because it will create multiple S3 bucket policies, which is not possible as only one bucket policy can be attached to an S3 bucket. Option B is not optimal because it will use the aws:SourceVpc condition key, which matches the ID of the VPC that is used to access the S3 bucket, not the VPC endpoint. Option C is not optimal because it will use the StringNotEquals condition operator with a single value, which will deny access to the S3 bucket from all VPC endpoints except one.
References: Using Amazon S3 Bucket Policies and User Policies, AWS Global Condition Context Keys

NEW QUESTION 80
A company is migrating its PostgreSQL database into the AWS Cloud. The company wants to use a database that will secure and regularly rotate database credentials. The company wants a solution that does not require additional programming overhead.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
This solution meets the requirements because it uses a PostgreSQL- compatible database that can secure and regularly rotate database credentials without requiring additional programming overhead. Amazon Aurora PostgreSQL is a relational database service that is compatible with PostgreSQL and offers high performance, availability, and scalability. AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. You can store database credentials in AWS Secrets Manager and use them to access your Aurora PostgreSQL database. You can also enable automatic rotation of your secrets according to a schedule or an event. AWS Secrets Manager handles the complexity of rotating secrets for you, such as generating new passwords and updating your database with the new credentials. Using Amazon DynamoDB for the database will not meet the requirements because it is a NoSQL database that is not compatible with PostgreSQL. Using AWS Systems Manager Parameter Store for storing and rotating database credentials will require additional programming overhead to integrate with your database.
Reference: [What Is Amazon Aurora?], [What Is AWS Secrets Manager?]

NEW QUESTION 82
A developer accesses AWS CodeCommit over SSH. The SSH keys configured to access AWS CodeCommit are tied to a user with the following permissions:

```
{
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "codecommit:BatchGetRepositories",
      "codecommit:Get*",
      "codecommit:List*",
      "codecommit:GitPull"
    ],
    "Resource": "*"
  }
]
}
```

The developer needs to create/delete branches
Which specific IAM permissions need to be added based on the principle of least privilege?

A. `"codecommit:CreateBranch"`
    `"codecommit:DeleteBranch"`

B. `"codecommit:Put*"`

C. `"codecommit:Update*"`

D. `"codecommit:*"`

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

Explanation:
This solution allows the developer to create and delete branches in AWS CodeCommit by granting the codecommit:CreateBranch and codecommit:DeleteBranch permissions. These are the minimum permissions required for this task, following the principle of least privilege. Option B grants too many permissions, such as codecommit:Put*, which allows the developer to create, update, or delete any resource in CodeCommit. Option C grants too few permissions, such as codecommit:Update*, which does not allow the developer to create or delete branches. Option D grants all permissions, such as codecommit:*, which is not secure or recommended.
Reference: [AWS CodeCommit Permissions Reference], [Create a Branch (AWS CLI)]

NEW QUESTION 87
A developer is troubleshooting an application mat uses Amazon DynamoDB in the uswest- 2 Region. The application is deployed to an Amazon EC2 instance. The application requires read-only permissions to a table that is named Cars The EC2 instance has an attached IAM role that contains the following IAM policy.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ReadOnlyAPIActions",
            "Effect": "Allow",
            "Action": [
                "dynamodb:GetItem",
                "dynamodb:BatchGetItem",
                "dynamodb:Scan",
                "dynamodb:Query",
                "dynamodb:ConditionCheckItem"
            ],
            "Resource": "arn:aws:dynamodb:us-west-2:account-id:table/Cars"
        }
    ]
}
```

When the application tries to read from the Cars table, an Access Denied error occurs. How can the developer resolve this error?

A. Modify the IAM policy resource to be "arn aws dynamo* us-west-2 account-id table/*"
B. Modify the IAM policy to include the dynamodb * action
C. Create a trust policy that specifies the EC2 service principa
D. Associate the role with the policy.
E. Create a trust relationship between the role and dynamodb Amazonas com.

Answer: C

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/access-control- overview.html#access-control-resource-ownership

NEW QUESTION 92
A developer is preparing to begin development of a new version of an application. The previous version of the application is deployed in a production environment. The developer needs to deploy fixes and updates to the current version during the development of the new version of the application. The code for the new version of the application is stored in AWS CodeCommit.
Which solution will meet these requirements?

A. From the main branch, create a feature branch for production bug fixe
B. Create a second feature branch from the main branch for development of the new version.
C. Create a Git tag of the code that is currently deployed in productio
D. Create a Git tag for the development of the new versio
E. Push the two tags to the CodeCommit repository.
F. From the main branch, create a branch of the code that is currently deployed in productio
G. Apply an IAM policy that ensures no other other users can push or merge to the branch.
H. Create a new CodeCommit repository for development of the new version of the applicatio
I. Create a Git tag for the development of the new version.

Answer: A

Explanation:
? A feature branch is a branch that is created from the main branch to work on a specific feature or task1. Feature branches allow developers to isolate their work from the main branch and avoid conflicts with other changes1. Feature branches can be merged back to the main branch when the feature or task is completed and tested1.
? In this scenario, the developer needs to maintain two parallel streams of work: one for fixing and updating the current version of the application that is deployed in production, and another for developing the new version of the application. The developer can use feature branches to achieve this goal.
? The developer can create a feature branch from the main branch for production bug fixes. This branch will contain the code that is currently deployed in production, and any fixes or updates that need to be applied to it. The developer can push this branch to the CodeCommit repository and use it to deploy changes to the production environment.
? The developer can also create a second feature branch from the main branch for development of the new version of the application. This branch will contain the code that is under development for the new version, and any changes or enhancements that are part of it. The developer can push this branch to the CodeCommit repository and use it to test and deploy the new version of the application in a separate environment.
? By using feature branches, the developer can keep the main branch stable and clean, and avoid mixing code from different versions of the application. The developer can also easily switch between branches and merge them when needed.

NEW QUESTION 96
A development team maintains a web application by using a single AWS CloudFormation template. The template defines web servers and an Amazon RDS database. The team uses the Cloud Formation template to deploy the Cloud Formation stack to different environments.
During a recent application deployment, a developer caused the primary development database to be dropped and recreated. The result of this incident was a loss of data. The team needs to avoid accidental database deletion in the future.
Which solutions will meet these requirements? (Choose two.)

A. Add a CloudFormation Deletion Policy attribute with the Retain value to the database resource.
B. Update the CloudFormation stack policy to prevent updates to the database.
C̲: Modify the database to use a Multi-AZ deployment.
D̲: Create a CloudFormation stack set for the web application and database deployments.
E. Add a Cloud Formation DeletionPolicy attribute with the Retain value to the stack.

Answer: AB

Explanation:
AWS CloudFormation is a service that enables developers to model and provision AWS resources using templates. The developer can add a CloudFormation Deletion Policy attribute with the Retain value to the database resource. This will prevent the database from being deleted when the stack is deleted or updated. The developer can also update the CloudFormation stack policy to prevent updates to the database. This will prevent accidental changes to the database configuration or properties.
References:
? [What Is AWS CloudFormation? - AWS CloudFormation]
? [DeletionPolicy Attribute - AWS CloudFormation]
? [Protecting Resources During Stack Updates - AWS CloudFormation]

NEW QUESTION 100
A company has a web application that runs on Amazon EC2 instances with a custom Amazon Machine Image (AMI) The company uses AWS CloudFormation to provision the application The application runs in the us-east-1 Region, and the company needs to deploy the application to the us-west-1 Region
An attempt to create the AWS CloudFormation stack in us-west-1 fails. An error message states that the AMI ID does not exist. A developer must resolve this error with a solution that uses the least amount of operational overhead
Which solution meets these requirements?

A. Change the AWS CloudFormation templates for us-east-1 and us-west-1 to use an AWS AM
B. Relaunch the stack for both Regions.
C. Copy the custom AMI from us-east-1 to us-west-1. Update the AWS CloudFormation template for us-west-1 to refer to AMI ID for the copied AMI Relaunch the stack
D. Build the custom AMI in us-west-1 Create a new AWS CloudFormation template to launch the stack in us-west-1 with the new AMI ID
E. Manually deploy the application outside AWS CloudFormation in us-west-1.

Answer: B

Explanation:
https://aws.amazon.com/blogs/aws/ec2-ami-copy-between-regions/

NEW QUESTION 102
A company has an application that stores data in Amazon RDS instances. The application periodically experiences surges of high traffic that cause performance problems.
During periods of peak traffic, a developer notices a reduction in query speed in all database queries.
The team's technical lead determines that a multi-threaded and scalable caching solution should be used to offload the heavy read traffic. The solution needs to improve performance.
Which solution will meet these requirements with the LEAST complexity?

A. Use Amazon ElastiCache for Memcached to offload read requests from the main database.
B. Replicate the data to Amazon DynamoD
C. Set up a DynamoDB Accelerator (DAX) cluster.
D. Configure the Amazon RDS instances to use Multi-AZ deployment with one standby instanc
E. Offload read requests from the main database to the standby instance.
F. Use Amazon ElastiCache for Redis to offload read requests from the main database.

Answer: A

Explanation:
? Amazon ElastiCache for Memcached is a fully managed, multithreaded, and scalable in-memory key-value store that can be used to cache frequently accessed

data and improve application performance1. By using Amazon ElastiCache for Memcached, the developer can reduce the load on the main database and handle high traffic surges more efficiently.
? To use Amazon ElastiCache for Memcached, the developer needs to create a cache cluster with one or more nodes, and configure the application to store and retrieve data from the cache cluster2. The developer can use any of the supported Memcached clients to interact with the cache cluster3. The developer can also use Auto Discovery to dynamically discover and connect to all cache nodes in a cluster4.
? Amazon ElastiCache for Memcached is compatible with the Memcached protocol, which means that the developer can use existing tools and libraries that work with
Memcached1. Amazon ElastiCache for Memcached also supports data partitioning, which allows the developer to distribute data among multiple nodes and scale out the cache cluster as needed.
? Using Amazon ElastiCache for Memcached is a simple and effective solution that meets the requirements with the least complexity. The developer does not need to change the database schema, migrate data to a different service, or use a different caching model. The developer can leverage the existing Memcached ecosystem and easily integrate it with the application.

NEW QUESTION 107
A developer has an application that stores data in an Amazon S3 bucket. The application uses an HTTP API to store and retrieve objects. When the PutObject API operation adds objects to the S3 bucket the developer must encrypt these objects at rest by using server- side encryption with Amazon S3 managed keys (SSE-S3).
Which solution will meet this requirement?

A. Create an AWS Key Management Service (AWS KMS) ke
B. Assign the KMS key to the S3 bucket.
C. Set the x-amz-server-side-encryption header when invoking the PutObject API operation.
D. Provide the encryption key in the HTTP header of every request.
E. Apply TLS to encrypt the traffic to the S3 bucket.

Answer: B

Explanation:
Amazon S3 supports server-side encryption, which encrypts data at rest on the server that stores the data. One of the encryption options is SSE-S3, which uses keys managed by S3. To use SSE-S3, the x-amz-server-side-encryption header must be set to AES256 when invoking the PutObject API operation. This instructs
downloaded. Reference:
S3 to encrypt the object data with SSE-S3 before saving it on disks in its data centers and decrypt it when it is
Protecting data using server-side encryption with Amazon S3-managed encryption keys (SSE-S3)

NEW QUESTION 110
A developer is trying get data from an Amazon DynamoDB table called demoman-table. The developer configured the AWS CLI to use a specific IAM use's credentials and ran the following command.

```
aws dynamodb get-item --table-name demoman-table --key '{"id": {"N":"1993"}}'
```

The command returned errors and no rows were returned. What is the MOST likely cause of these issues?

A. The command is incorrect; it should be rewritten to use put-item with a string argument
B. The developer needs to log a ticket with AWS Support to enable access to the demoman-table
C. Amazon DynamoOB cannot be accessed from the AWS CLI and needs to called via the REST API
D. The IAM user needs an associated policy with read access to demoman-table

Answer: D

Explanation:
This solution will most likely solve the issues because it will grant the IAM user the necessary permission to access the DynamoDB table using the AWS CLI command. The error message indicates that the IAM user does not have sufficient access rights to perform the scan operation on the table. Option A is not optimal because it will change the command to use put-item instead of scan, which will not achieve the desired result of getting data from the table. Option B is not optimal because it will involve contacting AWS Support, which may not be necessary or efficient for this issue. Option C is not optimal because it will state that DynamoDB cannot be accessed from the AWS CLI, which is incorrect as DynamoDB supports AWS CLI commands.
References: AWS CLI for DynamoDB, [IAM Policies for DynamoDB]

NEW QUESTION 113
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions.
When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

A. Set the Deployment Preference Type to Canary10Percent10Minute
AutoPublishAlias property to the Lambda alias.
B. Set the Deployment Preference Type to LinearlOPercentEvery10Minute
C. Set the
D. Set AutoPublishAlias property to the Lambda alias.
E. Set the Deployment Preference Type to CanaryIOPercentIOMinute
F. Set the PreTraffic and PostTraffic properties to the Lambda alias.
G. Set the Deployment Preference Type to LinearlOPercentEveryIOMinute
H. Set PreTraffic and Post Traffic properties to the Lambda alias.

Answer: A

Explanation:
The AWS Serverless Application Model (AWS SAM) comes built-in with CodeDeploy to provide gradual AWS Lambda deployments1.
The DeploymentPreference property in AWS SAM allows you to specify the type of deployment that you want. The Canary10Percent10Minutes option means that 10 percent of your customer traffic is immediately shifted to your new version. After 10 minutes, all traffic is shifted to the new version1. The AutoPublishAlias property in AWS SAM allows AWS SAM to automatically create an alias that points to the updated version of the Lambda function1. Therefore, option A is correct.

NEW QUESTION 114
A developer is writing an application that will retrieve sensitive data from a third-party system. The application will format the data into a PDF file. The PDF file could be more than 1 MB. The application will encrypt the data to disk by using AWS Key Management Service (AWS KMS). The application will decrypt the file when a user requests to download it. The retrieval and formatting portions of the application are complete.
The developer needs to use the GenerateDataKey API to encrypt the PDF file so that the PDF file can be decrypted later. The developer needs to use an AWS KMS symmetric customer managed key for encryption.
Which solutions will meet these requirements?

A. Write the encrypted key from the GenerateDataKey API to disk for later us
B. Use the plaintext key from the GenerateDataKey API and a symmetric encryption algorithm to encrypt the file.
C. Write the plain text key from the GenerateDataKey API to disk for later us
D. Use the encrypted key from the GenerateDataKey API and a symmetric encryption algorithm to encrypt the file.
E. Write the encrypted key from the GenerateDataKey API to disk for later us
F. Use the plaintext key from the GenerateDataKey API to encrypt the file by using the KMS Encrypt API
G. Write the plain text key from the GenerateDataKey API to disk for later us
H. Use the encrypted key from the GenerateDataKey API to encrypt the file by using the KMS Encrypt API

Answer: A

Explanation:
? The GenerateDataKey API returns a data key that is encrypted under a symmetric encryption KMS key that you specify, and a plaintext copy of the same data key1. The data key is a random byte string that can be used with any standard encryption algorithm, such as AES or SM42. The plaintext data key can be used to encrypt or decrypt data outside of AWS KMS, while the encrypted data key can be stored with the encrypted data and later decrypted by AWS KMS1.
? In this scenario, the developer needs to use the GenerateDataKey API to encrypt
the PDF file so that it can be decrypted later. The developer also needs to use an AWS KMS symmetric customer managed key for encryption. To achieve this, the developer can follow these steps:

NEW QUESTION 116
An application that runs on AWS Lambda requires access to specific highly confidential objects in an Amazon S3 bucket. In accordance with the principle of least privilege a company grants access to the S3 bucket by using only temporary credentials.
How can a developer configure access to the S3 bucket in the MOST secure way?

A. Hardcode the credentials that are required to access the S3 objects in the application cod
B. Use the credentials to access me required S3 objects.
C. Create a secret access key and access key ID with permission to access the S3 bucke
D. Store the key and key ID in AWS Secrets Manage
E. Configure the application to retrieve the Secrets Manager secret and use the credentials to access me S3 objects.
F. Create a Lambda function execution role Attach a policy to the rote that grants access to specific objects in the S3 bucket.
G. Create a secret access key and access key ID with permission to access the S3 bucket Store the key and key ID as environment variables m Lambd
H. Use the environment variables to access the required S3 objects.

Answer: C

Explanation:
This solution will meet the requirements by creating a Lambda function execution role, which is an IAM role that grants permissions to a Lambda function to access AWS resources such as Amazon S3 objects. The developer can attach a policy to the role that grants access to specific objects in the S3 bucket that are required by the application, following the principle of least privilege. Option A is not optimal because it will hardcode the credentials that are required to access S3 objects in the application code, which is insecure and difficult to maintain. Option B is not optimal because it will create a secret access key and access key ID with permission to access the S3 bucket, which will introduce additional security risks and complexity for storing and managing credentials. Option D is not optimal because it will store the secret access key and access key ID as environment variables in Lambda, which is also insecure and difficult to maintain. References: [AWS Lambda Execution Role], [Using AWS Lambda with Amazon S3]

NEW QUESTION 117
An organization is using Amazon CloudFront to ensure that its users experience low- latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application.
How can these requirements be met? (Select TWO)

A. Use AWS KMS t0 encrypt traffic between cloudFront and the web application.
B. Set the Origin Protocol Policy to "HTTPS Only".
C. Set the Origin's HTTP Port to 443.
D. Set the Viewer Protocol Policy to "HTTPS Only" or Redirect HTTP to HTTPS"
E. Enable the CloudFront option Restrict Viewer Access.

Answer: BD

Explanation:
This solution will meet the requirements by ensuring that all traffic between users and CloudFront, and all traffic between CloudFront and the web application, are encrypted using HTTPS protocol. The Origin Protocol Policy determines how CloudFront communicates with the origin server (the web application), and setting it to "HTTPS Only" will force CloudFront to use HTTPS for every request to the origin server. The Viewer Protocol Policy determines how CloudFront responds to HTTP or HTTPS requests from users, and setting it to "HTTPS Only" or "Redirect HTTP to HTTPS" will force CloudFront to use HTTPS for every response to users. Option A is not optimal because it will use AWS KMS to encrypt traffic between CloudFront and the web application, which is not necessary or supported by CloudFront. Option C is not optimal because it will set the origin's HTTP port to 443, which is incorrect as port 443 is used for HTTPS protocol, not HTTP protocol. Option E is not optimal because it will enable the CloudFront option Restrict Viewer Access, which is used for controlling access to private content using signed URLs or signed cookies, not for encrypting traffic.
References: [Using HTTPS with CloudFront], [Restricting Access to Amazon S3 Content by Using an Origin Access Identity]

NEW QUESTION 122
A company is running Amazon EC2 instances in multiple AWS accounts. A developer needs to implement an application that collects all the lifecycle events of the EC2 instances. The application needs to store the lifecycle events in a single Amazon Simple Queue Service (Amazon SQS) queue in the company's main AWS account for further processing.

Which solution will meet these requirements?

A. Configure Amazon EC2 to deliver the EC2 instance lifecycle events from all accounts to the Amazon EventBridge event bus of the main accoun
B. Add an EventBridge rule to the event bus of the main account that matches all EC2 instance lifecycle event
C. Add the SQS queue as a target of the rule.
D. Use the resource policies of the SQS queue in the main account to give each account permissions to write to that SQS queu
E. Add to the Amazon EventBridge event bus of each account an EventBridge rule that matches all EC2 instance lifecycle event
F. Add the SQS queue in the main account as a target of the rule.
G. Write an AWS Lambda function that scans through all EC2 instances in the company accounts to detect EC2 instance lifecycle change
H. Configure the Lambda function to write a notification message to the SQS queue in the main account if the function detects an EC2 instance lifecycle chang
I. Add an Amazon EventBridge scheduled rule that invokes the Lambda function every minute.
J. Configure the permissions on the main account event bus to receive events from all account
K. Create an Amazon EventBridge rule in each account to send all the EC2 instance lifecycle events to the main account event bu
L. Add an EventBridge rule to the main account event bus that matches all EC2 instance lifecycle event
M. Set the SQS queue as a target for the rule.

Answer: D

Explanation:
 Amazon EC2 instances can send the state-change notification events to Amazon EventBridge.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-instance-state- changes.html Amazon EventBridge can send and receive events between event buses in AWS accounts. https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross- account.html

NEW QUESTION 123
A company has an analytics application that uses an AWS Lambda function to process transaction data asynchronously A developer notices that asynchronous invocations of the Lambda function sometimes fail When failed Lambda function invocations occur, the developer wants to invoke a second Lambda function to handle errors and log details.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
 Configuring a Lambda function destination with a failure condition is the best solution for invoking a second Lambda function to handle errors and log details. A Lambda function destination is a resource that Lambda sends events to after a function is invoked. The developer can specify the destination type as Lambda function and the ARN of the error-handling Lambda function as the resource. The developer can also specify the failure condition, which means that the destination is invoked only when the initial Lambda function fails. The destination event will include the response from the initial function, the request ID, and the timestamp. The other solutions are either not feasible or not efficient. Enabling AWS X-Ray active tracing on the initial Lambda function will help to monitor and troubleshoot the function performance, but it will not automatically invoke the error-handling Lambda function. Configuring a Lambda function trigger with a failure condition is not a valid option, as triggers are used to invoke Lambda functions, not to send events from Lambda functions. Creating a status check alarm on the initial Lambda function will incur additional costs and complexity, and it will not capture the details of the failed
invocations. References
? Using AWS Lambda destinations
? Asynchronous invocation - AWS Lambda
? AWS Lambda Destinations: What They Are and Why to Use Them
? AWS Lambda Destinations: A Complete Guide | Dashbird

NEW QUESTION 126
A developer has an application that is composed of many different AWS Lambda functions. The Lambda functions all use some of the same dependencies. To avoid security issues the developer is constantly updating the dependencies of all of the Lambda functions. The result is duplicated effort to reach function.
How can the developer keep the dependencies of the Lambda functions up to date with the LEAST additional complexity?

A. Define a maintenance window for the Lambda functions to ensure that the functions get updated copies of the dependencies.
B. Upgrade the Lambda functions to the most recent runtime version.
C. Define a Lambda layer that contains all of the shared dependencies.
D. Use an AWS CodeCommit repository to host the dependencies in a centralized location.

Answer: C

Explanation:
 This solution allows the developer to keep the dependencies of the Lambda functions up to date with the least additional complexity because it eliminates the need to update each function individually. A Lambda layer is a ZIP archive that contains libraries, custom runtimes, or other dependencies. The developer can create a layer that contains all of the shared dependencies and attach it to multiple Lambda functions. When the developer updates the layer, all of the functions that use the layer will have access to the latest version of the dependencies.
Reference: [AWS Lambda layers]

NEW QUESTION 131
A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values
When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.
Which solution will meet these requirements with the LEAST development effort?

A. Modify the CloudFormation stack to set the deletion policy to Retain for the Parameter Store parameters.
B. Create an Amazon DynamoDB table as a resource in the CloudFormation stack to hold configuration data for the application Migrate the parameters that the application is modifying from Parameter Store to the DynamoDB table

C. Create an Amazon RDS DB instance as a resource in the CloudFormation stac
D. Create a table in the database for parameter configuratio
E. Migrate the parameters that the application is modifying from Parameter Store to the configuration table
F. Modify the CloudFormation stack policy to deny updates on Parameter Store parameters

Answer: D

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack- resources.html#stack-policy-samples

NEW QUESTION 134
A developer created an AWS Lambda function that performs a series of operations that involve multiple AWS services. The function's duration time is higher than normal. To determine the cause of the issue, the developer must investigate traffic between the services without changing the function code
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
 AWS X-Ray is a service that helps you analyze and debug your applications. You can use X-Ray to trace requests made to your Lambda function and other AWS services, and identify performance bottlenecks and errors. Enabling active tracing in your Lambda function allows X-Ray to collect data from the function invocation and the downstream services that it calls. You can then review the logs and service maps in X-Ray to diagnose the issue. References
? Monitoring and troubleshooting Lambda functions - AWS Lambda
? Using AWS Lambda with AWS X-Ray
? Troubleshoot Lambda function cold start issues | AWS re:Post

NEW QUESTION 136
A developer is using AWS Amplify Hosting to build and deploy an application. The developer is receiving an increased number of bug reports from users. The developer wants to add end-to-end testing to the application to eliminate as many bugs as possible before the bugs reach production.
Which solution should the developer implement to meet these requirements?

A. Run the amplify add test command in the Amplify CLI.
B. Create unit tests in the applicatio
C. Deploy the unit tests by using the amplify push command in the Amplify CLI.
D. Add a test phase to the amplify.yml build settings for the application.
E. Add a test phase to the aws-exports.js file for the application.

Answer: C

Explanation:
 The solution that will meet the requirements is to add a test phase to the amplify.yml build settings for the application. This way, the developer can run end-to-end tests on every code commit and catch any bugs before deploying to production. The other options either do not support end-to-end testing, or do not run tests automatically.
Reference: End-to-end testing

NEW QUESTION 140
A company has deployed infrastructure on AWS. A development team wants to create an AWS Lambda function that will retrieve data from an Amazon Aurora database. The Amazon Aurora database is in a private subnet in company's VPC. The VPC is named VPC1. The data is relational in nature. The Lambda function needs to access the data
                              securely.
Which solution will meet these requirements?

A. Create the Lambda functio
B. Configure VPC1 access for the functio
C. Attach a security group named SG1 to both the Lambda function and the databas
D. Configure the security group inbound and outbound rules to allow TCP traffic on Port 3306.
E. Create and launch a Lambda function in a new public subnet that is in a new VPC named VPC2. Create a peering connection between VPC1 and VPC2.
F. Create the Lambda functio
G. Configure VPC1 access for the functio
H. Assign a security group named SG1 to the Lambda functio
I. Assign a second security group named SG2 to the databas
J. Add an inbound rule to SG1 to allow TCP traffic from Port 3306.
K. Export the data from the Aurora database to Amazon S3. Create and launch a Lambda function in VPC1. Configure the Lambda function query the data from Amazon S3.

Answer: A

Explanation:
 AWS Lambda is a service that lets you run code without provisioning or managing servers. Lambda functions can be configured to access resources in a VPC, such as an Aurora database, by specifying one or more subnets and security groups in the VPC settings of the function. A security group acts as a virtual firewall that controls inbound and outbound traffic for the resources in a VPC. To allow a Lambda function to communicate with an Aurora database, both resources need to be associated with the same security group, and the security group rules need to allow TCP traffic on Port 3306, which is the default port for MySQL databases.
Reference: [Configuring a Lambda function to access resources in a VPC]

NEW QUESTION 144

A developer has a legacy application that is hosted on-premises. Other applications hosted on AWS depend on the on-premises application for proper functioning. In case of any application errors, the developer wants to be able to use Amazon CloudWatch to monitor and troubleshoot all applications from one place. How can the developer accomplish this?

A. Install an AWS SDK on the on-premises server to automatically send logs to CloudWatch.
B. Download the CloudWatch agent to the on-premises serve
C. Configure the agent to use IAM user credentials with permissions for CloudWatch.
D. Upload log files from the on-premises server to Amazon S3 and have CloudWatch read the files.
E. Upload log files from the on-premises server to an Amazon EC2 instance and have the instance forward the logs to CloudWatch.

Answer: B

Explanation:
 Amazon CloudWatch is a service that monitors AWS resources and applications. The developer can use CloudWatch to monitor and troubleshoot all applications from one place. To do so, the developer needs to download the CloudWatch agent to the on-premises server and configure the agent to use IAM user credentials with permissions for CloudWatch. The agent will collect logs and metrics from the on-premises server and send them to CloudWatch.
References:
                                    ? [What Is Amazon CloudWatch? - Amazon CloudWatch]
? [Installing and Configuring the CloudWatch Agent - Amazon CloudWatch]

NEW QUESTION 146
A developer deployed an application to an Amazon EC2 instance The application needs to know the public IPv4 address of the instance
How can the application find this information?

                                    Query the instance metadata from http./M69.254.169.254. latestmeta-data/.
A.
B: Query the instance user data from http '169 254.169 254. latest/user-data/
C. Query the Amazon Machine Image (AMI) information from http://169.254.169.254/latest/meta-data/ami/.
D. Check the hosts file of the operating system

Answer: A

Explanation:
 The instance metadata service provides information about the EC2 instance, including the public IPv4 address, which can be obtained by querying the endpoint http://169.254.169.254/latest/meta-data/public-ipv4. References
? Instance metadata and user data
? Get Public IP Address on current EC2 Instance
? Get the public ip address of your EC2 instance quickly

NEW QUESTION 149
A team of developed is using an AWS CodePipeline pipeline as a continuous integration and continuous delivery (CI/CD) mechanism for a web application. A developer has written unit tests to programmatically test the functionality of the application code. The unit tests produce a test report that shows the results of each individual check. The developer now
                                    wants to run these tests automatically during the CI/CD process.

A. Write a Git pre-commit hook that runs the test before every commi
B. Ensure that each developer who is working on the project has the pre-commit hook instated locall
C. Review the test report and resolve any issues before pushing changes to AWS CodeCommit.
D. Add a new stage to the pipelin
E. Use AWS CodeBuild as the provide
F. Add the new stage after the stage that deploys code revisions to the test environmen
G. Write a buildspec that fails the CodeBuild stage if any test does not pas
H. Use the test reports feature of Codebuild to integrate the report with the CodoBuild consol
I. View the test results in CodeBuild Resolve any issues.
J. Add a new stage to the pipelin
K. Use AWS CodeBuild at the provide
L. Add the new stage before the stage that deploys code revisions to the test environmen
M. Write a buildspec that fails the CodeBuild stage it any test does not pas
N. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild consol
O. View the test results in codeBuild Resolve any issues.
P. Add a new stage to the pipelin
Q. Use Jenkins as the provide
R. Configure CodePipeline to use Jenkins to run the unit test
S. Write a Jenkinsfile that fails the stage if any test does not pas
T. Use the test report plugin for Jenkins to integrate the repot with the Jenkins dashboar
. View the test results in Jenkin
. Resolve any issues.

Answer: C

Explanation:
 The solution that will meet the requirements is to add a new stage to the pipeline. Use AWS CodeBuild as the provider. Add the new stage before the stage that deploys code revisions to the test environment. Write a buildspec that fails the CodeBuild stage if any test does not pass. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild console. View the test results in CodeBuild. Resolve any issues. This way, the developer can run the unit tests automatically during the CI/CD process and catch any bugs before deploying to the test environment. The developer can also use the test reports feature of CodeBuild to view and analyze the test results in a graphical interface. The other options either involve running the tests manually, running them after deployment, or using a different provider that requires additional configuration and integration.
Reference: Test reports for CodeBuild

NEW QUESTION 152

A company's website runs on an Amazon EC2 instance and uses Auto Scaling to scale the environment during peak times. Website users across the world ate experiencing high latency flue lo sialic content on theEC2 instance. even during non-peak hours.
When companion of steps mill resolves the latency issue? (Select TWO)

A. Double the Auto Scaling group's maximum number of servers
B. Host the application code on AWS lambda
C. Scale vertically by resizing the EC2 instances
D. Create an Amazon Cloudfront distribution to cache the static content
E. Store the application's sialic content in Amazon S3

Answer: DE

Explanation:
 The combination of steps that will resolve the latency issue is to create an Amazon CloudFront distribution to cache the static content and store the application's static content in Amazon S3. This way, the company can use CloudFront to deliver the static content from edge locations that are closer to the website users, reducing latency and improving performance. The company can also use S3 to store the static content reliably and cost-effectively, and integrate it with CloudFront easily. The other options either do not address the latency issue, or are not necessary or feasible for the given scenario.
Reference: Using Amazon S3 Origins and Custom Origins for Web Distributions

NEW QUESTION 153
A company runs a batch processing application by using AWS Lambda functions and Amazon API Gateway APIs with deployment stages for development, user acceptance testing and production A development team needs to configure the APIs in the deployment stages to connect to third-party service endpoints.
Which solution will meet this requirement?

A. Store the third-party service endpoints in Lambda layers that correspond to the stage
B. Store the third-party service endpoints in API Gateway stage variables that correspond to the stage
C. Encode the third-party service endpoints as query parameters in the API Gateway request URL.
D. Store the third-party service endpoint for each environment in AWS AppConfig

Answer: B

Explanation:
 API Gateway stage variables are name-value pairs that can be defined as configuration attributes associated with a deployment stage of a REST API. They act like environment variables and can be used in the API setup and mapping templates. For example, the development team can define a stage variable named endpoint and assign it different values for each stage, such as dev.example.com for development, uat.example.com for user acceptance testing, and prod.example.com for production. Then, the team can use the stage variable value in the integration request URL, such as http://$ { stageVariables.endpoint}/api. This way, the team can use the same API setup with different endpoints at each stage by resetting the stage variable value. The other solutions are either not feasible or not cost-effective. Lambda layers are used to package and load dependencies for Lambda functions, not for storing endpoints. Encoding the endpoints as query parameters would expose them to the public and make the request URL unnecessarily long. Storing the endpoints in AWS AppConfig would incur additional costs   and complexity, and would require additional logic to retrieve the values from the configuration store. References
? Using Amazon API Gateway stage variables
? Setting up stage variables for a REST API deployment
? Setting stage variables using the Amazon API Gateway console

NEW QUESTION 157
A developer warns to add request validation to a production environment Amazon API Gateway API. The developer needs to test the changes before the API is deployed to the production environment. For the lest the developer will send test requests to the API through a testing tool.
Which solution will meet these requirements with the LEAST operational overhead?

A. Export the existing API to an OpenAPI fil
B. Create a new API Import the OpenAPI file Modify the new API to add request validatio
C. Perform the tests Modify the existing API to add request validatio
D. Deploy the existing API to production.
E. Modify the existing API to add request validatio
F. Deploy the updated API to a new API Gateway stage Perform the tests Deploy the updated API to the API Gateway production stage.
G. Create a new API Add the necessary resources and methods including new request validatio
H. Perform the tests Modify the existing API to add request validatio
I. Deploy the existing API to production.
J. Clone the exiting API Modify the new API lo add request validatio
                                                       Modify the existing API to add request validation Deploy the existing API to production.
K. Perform the tests

Answer: D

Explanation:
 This solution allows the developer to test the changes without affecting the production environment. Cloning an API creates a copy of the API definition that can be modified independently. The developer can then add request validation to the new API and test it using a testing tool. After verifying that the changes work as expected, the developer can apply the same changes to the existing API and deploy it to production.
Reference: Clone an API, [Enable Request Validation for an API in API Gateway]

NEW QUESTION 161
A developer migrated a legacy application to an AWS Lambda function. The function uses a third-party service to pull data with a series of API calls at the end of each month. The function than processes the data to generate the monthly reports. The function has Been working with no issues so far.
The third-party service recently issued a restriction to allow a feed number to API calls each minute and each day. If the API calls exceed the limit tor each minute or each day, then the service will produce errors. The API also provides the minute limit and daily limit in the response header. This restriction might extend the overall process to multiple days because the process is consuming more API calls than the available limit.
What is the MOST operationally efficient way to refactor the server less application to accommodate this change?

A. Use an AWS Step Functions State machine to monitor API failure
B. Use the Wait state to delay calling the Lambda function.

C. Use an Amazon Simple Queue Service (Amazon SQS) queue to hold the API call
D. Configure the Lambda function to poll the queue within the API threshold limits.
Use an Amazon CloudWatch Logs metric to count the number of API call
F: Configure an Amazon CloudWatch alarm flat slops the currently running instance of the Lambda function when the metric exceeds the API threshold limits.
G. Use Amazon Kinesis Data Firehose to batch me API calls and deliver them to an Amazon S3 bucket win an event notification to invoke the Lambda function.

Answer: A

Explanation:
 The solution that will meet the requirements is to use an AWS Step Functions state machine to monitor API failures. Use the Wait state to delay calling the Lambda function. This way, the developer can refactor the serverless application to accommodate the change in a way that is automated and scalable. The developer can use Step Functions to orchestrate the Lambda function and handle any errors or retries. The developer can also use the Wait state to pause the execution for a specified duration or until a specified timestamp, which can help avoid exceeding the API limits. The other options either involve using additional services that are not necessary or appropriate for this scenario, or do not address the issue of API failures.
Reference: AWS Step Functions Wait state

NEW QUESTION 162
A company is using Amazon RDS as the Backend database for its application. After a recent marketing campaign, a surge of read requests to the database increased the latency of data retrieval from the database.
The company has decided to implement a caching layer in front of the database. The cached content must be encrypted and must be highly available.
Which solution will meet these requirements?

A. Amazon Cloudfront
B. Amazon ElastiCache to Memcached
C. Amazon ElastiCache for Redis in cluster mode
D. Amazon DynamoDB Accelerate (DAX)

Answer: C

Explanation:
 This solution meets the requirements because it provides a caching layer that can store and retrieve encrypted data from multiple nodes. Amazon ElastiCache for Redis supports encryption at rest and in transit, and can scale horizontally to increase the cache capacity and availability. Amazon ElastiCache for Memcached does not support encryption, Amazon CloudFront is a content delivery network that is not suitable for caching database queries, and Amazon DynamoDB Accelerator (DAX) is a caching service that only works with DynamoDB tables.
Reference: [Amazon ElastiCache for Redis Features], [Choosing a Cluster Engine]

NEW QUESTION 167
A developer is creating an AWS Lambda function in VPC mode An Amazon S3 event will invoke the Lambda function when an object is uploaded into an S3 bucket The Lambda function will process the object and produce some analytic results that will be recorded into a file Each processed object will also generate a log entry that will be recorded into a file.
Other Lambda functions. AWS services, and on-premises resources must have access to the result files and log file. Each log entry must also be appended to the same shared log file. The developer needs a solution that can share files and append results into an existing file.
Which solution should the developer use to meet these requirements?

A. Create an Amazon Elastic File System (Amazon EFS) file syste
B. Mount the EFS file system in Lambd
C. Store the result files and log file in the mount poin
D. Append the log entries to the log file.
E. Create an Amazon Elastic Block Store (Amazon EBS) Multi-Attach enabled volume Attach the EBS volume to all Lambda function
  download the log file, append the log entries, and upload the modified log file to Amazon EBS
F. Update the Lambda function code to
G. Create a reference to the /tmp local director
H. Store the result files and log file by using the directory referenc
I. Append the log entry to the log file.
J. Create a reference to the /opt storage directory Store the result files and log file by using the directory reference Append the log entry to the log file

Answer: A

Explanation:
 https://aws.amazon.com/blogs/compute/using-amazon-efs-for-aws-lambda-in-your-serverless-applications/

NEW QUESTION 168
A developer is building a web application that uses Amazon API Gateway to expose an AWS Lambda function to process requests from clients. During testing, the developer notices that the API Gateway times out even though the Lambda function finishes under the set time limit.
Which of the following API Gateway metrics in Amazon CloudWatch can help the developer troubleshoot the issue? (Choose two.)

A. CacheHitCount
B. IntegrationLatency
C. CacheMissCount
D. Latency
E. Count

Answer: BD

Explanation:
 Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. Amazon CloudWatch is a service that monitors AWS resources and applications. API Gateway provides several CloudWatch metrics to help developers troubleshoot issues with their APIs. Two of the metrics that can help the developer troubleshoot the issue of API Gateway timing out are:
? IntegrationLatency: This metric measures the time between when API Gateway
relays a request to the backend and when it receives a response from the backend. A high value for this metric indicates that the backend is taking too long to

respond and may cause API Gateway to time out.
? Latency: This metric measures the time between when API Gateway receives a
request from a client and when it returns a response to the client. A high value for this metric indicates that either the integration latency is high or API Gateway is taking too long to process the request or response.
References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [Amazon API Gateway Metrics and Dimensions - Amazon CloudWatch]
? [Troubleshooting API Errors - Amazon API Gateway]

NEW QUESTION 171
A developer is designing a serverless application with two AWS Lambda functions to process photos. One Lambda function stores objects in an Amazon S3 bucket and stores the associated metadata in an Amazon DynamoDB table. The other Lambda function fetches the objects from the S3 bucket by using the metadata from the DynamoDB table. Both Lambda functions use the same Python library to perform complex computations and are approaching the quota for the maximum size of zipped deployment packages.
What should the developer do to reduce the size of the Lambda deployment packages with the LEAST operational overhead?

A. Package each Python library in its own .zip file archiv
B. Deploy each Lambda function with its own copy of the library.
C. Create a Lambda layer with the required Python librar
D. Use the Lambda layer in both Lambda functions.
E. Combine the two Lambda functions into one Lambda functio
F. Deploy the Lambda function as a single .zip file archive.
G. Download the Python library to an S3 bucke
H. Program the Lambda functions to reference the object URLs.

Answer: B

Explanation:
AWS Lambda is a service that lets developers run code without provisioning or managing servers. Lambda layers are a distribution mechanism for libraries, custom runtimes, and other dependencies. The developer can create a Lambda layer with the

required Python library and use the layer in both Lambda functions. This will reduce the size of the Lambda deployment packages and avoid reaching the quota for the maximum size of zipped deployment packages. The developer can also benefit from using layers to manage dependencies separately from function code.
References:
? [What Is AWS Lambda? - AWS Lambda]
? [AWS Lambda Layers - AWS Lambda]

NEW QUESTION 174
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment, the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

A.

All at once
B. Rolling with additional batch
C. Blue/green
D. Immutable

Answer: D

Explanation:
The immutable deployment method is the best option for this scenario, because it meets the requirements of maintaining full capacity, avoiding service interruption, and minimizing the cost of additional resources.
The immutable deployment method creates a new set of instances in a separate Auto Scaling group and deploys the new version of the application to them. Then, it swaps the new instances with the old ones and terminates the old instances. This way, the application maintains full capacity during the deployment and avoids any downtime. The cost of additional resources is also minimized, because the new instances are only created for a short time and then replaced by the old ones.
The other deployment methods do not meet all the requirements:
? The all at once method deploys the new version to all instances simultaneously, which causes a short period of downtime and reduced capacity.
? The rolling with additional batch method deploys the new version in batches, but for the first batch it creates new instances instead of using the existing ones. This increases the cost of additional resources and reduces the capacity of the original environment.
? The blue/green method creates a new environment with a new set of instances and deploys the new version to them. Then, it swaps the URLs between the old and new environments. This method maintains full capacity and avoids service interruption, but it also increases the cost of additional resources significantly, because it duplicates the entire environment.

NEW QUESTION 179

A developer is creating a serverless application that uses an AWS Lambda function The developer will use AWS CloudFormation to deploy the application The application will write logs to Amazon CloudWatch Logs The developer has created a log group in a CloudFormation template for the application to use The developer needs to modify the CloudFormation template to make the name of the log group available to the application at runtime Which solution will meet this requirement?

A. Use the AWS:Include transform in CloudFormation to provide the log group's name to the application
B. Pass the log group's name to the application in the user data section of the CloudFormation template.
C. Use the CloudFormation template's Mappings section to specify the log group's name for the application.
D. Pass the log group's Amazon Resource Name (ARN) as an environment variable to the Lambda function

Answer: D

Explanation:
 FunctionName: MyLambdaFunction Code:
S3Bucket: your-lambda-code-bucket S3Key: lambda-code.zip
Runtime: nodejs14.x # Specify the desired runtime for your Lambda function Environment:
Variables:
LOG_GROUP_NAME: !Ref MyLogGroup https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-logs- loggroup.html

NEW QUESTION 184
A developer is troubleshooting an application in an integration environment. In the application, an Amazon Simple Queue Service (Amazon SQS) queue consumes messages and then an AWS Lambda function processes the messages. The Lambda function transforms the messages and makes an API call to a third-party service.
There has been an increase in application usage. The third-party API frequently returns an HTTP 429 Too Many Requests error message. The error message prevents a significant number of messages from being processed successfully.
How can the developer resolve this issue?

A. Increase the SQS event source's batch size setting.
B. Configure provisioned concurrency for the Lambda function based on the third-party API's documented rate limits.
C. Increase the retry attempts and maximum event age in the Lambda function's asynchronous configuration.
D. Configure maximum concurrency on the SQS event source based on the third-party service's documented rate limits.

Answer: D

Explanation:
? Maximum concurrency for SQS as an event source allows customers to control the maximum concurrent invokes by the SQS event source1. When multiple SQS event sources are configured to a function, customers can control the maximum concurrent invokes of individual SQS event source1.
? In this scenario, the developer needs to resolve the issue of the third-party API frequently returning an HTTP 429 Too Many Requests error message, which prevents a significant number of messages from being processed successfully. To achieve this, the developer can follow these steps:
? By using this solution, the developer can reduce the frequency of HTTP 429 errors and improve the message processing success rate. The developer can also avoid throttling or blocking by the third-party API.

NEW QUESTION 187
An ecommerce application is running behind an Application Load Balancer. A developer observes some unexpected load on the application during non-peak hours. The developer wants to analyze patterns for the client IP addresses that use the application. Which HTTP header should the developer use for this analysis?

A. The X-Forwarded-Proto header
B. The X-F Forwarded-Host header
C. The X-Forwarded-For header
D. The X-Forwarded-Port header

Answer: C

Explanation:
 The HTTP header that the developer should use for this analysis is the X- Forwarded-For header. This header contains the IP address of the client that made the request to the Application Load Balancer. The developer can use this header to analyze patterns for the client IP addresses that use the application. The other headers either contain information about the protocol, host, or port of the request, which are not relevant for the analysis.
Reference: How Application Load Balancer works with your applications

NEW QUESTION 190
A developer is creating a template that uses AWS CloudFormation to deploy an application. The application is serverless and uses Amazon API Gateway, Amazon DynamoDB, and AWS Lambda.
Which AWS service or tool should the developer use to define serverless resources in YAML?

A. CloudFormation serverless intrinsic functions

B. AWS Elastic Beanstalk
C. AWS Serverless Application Model (AWS SAM)
D. AWS Cloud Development Kit (AWS CDK)

Answer: C

Explanation:
 AWS Serverless Application Model (AWS SAM) is an open-source framework that enables developers to build and deploy serverless applications on AWS. AWS SAM uses a template specification that extends AWS CloudFormation to simplify the

definition of serverless resources such as API Gateway, DynamoDB, and Lambda. The developer can use AWS SAM to define serverless resources in YAML and deploy them using the AWS SAM CLI.
References:
? [What Is the AWS Serverless Application Model (AWS SAM)? - AWS Serverless Application Model]
? [AWS SAM Template Specification - AWS Serverless Application Model]

NEW QUESTION 193
A data visualization company wants to strengthen the security of its core applications The applications are deployed on AWS across its development staging, pre-production, and production environments. The company needs to encrypt all of its stored sensitive credentials The sensitive credentials need to be automatically rotated Aversion of the sensitive credentials need to be stored for each environment
Which solution will meet these requirements in the MOST operationally efficient way?

A. Configure AWS Secrets Manager versions to store different copies of the same credentials across multiple environments
B. Create a new parameter version in AWS Systems Manager Parameter Store for each environment Store the environment-specific credentials in the parameter version.
C. Configure the environment variables in the application code Use different names for each environment type
D. Configure AWS Secrets Manager to create a new secret for each environment type. Store the environment-specific credentials in the secret

Answer: D

Explanation:
 AWS Secrets Manager is the best option for managing sensitive credentials across multiple environments, as it provides automatic secret rotation, auditing, and monitoring features. It also allows storing environment-specific credentials in separate secrets, which can be accessed by the applications using the SDK or CLI. AWS Systems Manager Parameter Store does not have built-in secret rotation capability, and it requires creating individual parameters or storing the entire credential set as JSON object. Configuring the environment variables in the application code is not a secure or scalable solution, as it exposes the credentials to anyone who can access the code. References
? AWS Secrets Manager vs. Systems Manager Parameter Store
? AWS System Manager Parameter Store vs Secrets Manager vs Environment Variation in Lambda, when to use which
? AWS Secrets Manager vs. Parameter Store: Features, Cost & More

NEW QUESTION 194
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AWS-Certified-Developer-Associate Practice Exam Features:

* AWS-Certified-Developer-Associate Questions and Answers Updated Frequently

* AWS-Certified-Developer-Associate Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Developer-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Developer-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Developer-Associate Practice Test Here](https://www.certshared.com/exam/AWS-Certified-Developer-Associate/)