Number: DVA-C02
Passing Score: 800
Time Limit: 120
File Version: 13.0

Exam Code: DVA-C02

Exam Name: AWS Certified Developer - Associate

Exam A

QUESTION 1
A developer needs to store configuration variables for an application. The developer needs to set an expiration date and time for me configuration. The developer wants to receive notifications. Before the configuration expires. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a standard parameter in AWS Systems Manager Parameter Store Set Expiation and Expiration Notification policy types.
B. Create a standard parameter in AWS Systems Manager Parameter Store Create an AWS Lambda function to expire the configuration and to send Amazon Simple Notification Service (Amazon SNS) notifications.
C. Create an advanced parameter in AWS Systems Manager Parameter Store Set Expiration and Expiration Notification policy types.
D. Create an advanced parameter in AWS Systems Manager Parameter Store Create an Amazon EC2 instance with a corn job to expire the configuration and to send notifications.

Correct Answer: C
Section:
Explanation:
This solution will meet the requirements by creating an advanced parameter in AWS Systems Manager Parameter Store, which is a secure and scalable service for storing and managing configuration data and secrets. The advanced parameter allows setting expiration and expiration notification policy types, which enable specifying an expiration date and time for the configuration and receiving notifications before the configuration expires. 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 create a standard parameter in AWS Systems Manager Parameter Store, which does not support expiration and expiration notification policy types. 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 create a Docker container from Node.js base image to invoke Lambda functions, which will incur additional costs and overhead for creating and running Docker containers.
Reference: AWS Systems Manager Parameter Store, [Using SSL/TLS to Encrypt a Connection to a DB Instance]

QUESTION 2
When using the AWS Encryption SDK how does the developer keep track of the data encryption keys used to encrypt data?

A. The developer must manually keep Hack of the data encryption keys used for each data object.
B. The SDK encrypts me data encryption key and stores it (encrypted) as part of the resumed ophertext.
C. The SDK stores the data encryption keys automaticity in Amazon S3.
D. The data encryption key is stored m the user data for the EC2 instance.

Correct Answer: B
Section:
Explanation:
This solution will meet the requirements by using AWS Encryption SDK, which is a client-side encryption library that enables developers to encrypt and decrypt data using data encryption keys that are protected by AWS Key Management Service (AWS KMS). The SDK encrypts the data encryption key with a customer master key (CMK) that is managed by AWS KMS, and stores it (encrypted) as part of the returned ciphertext. The developer does not need to keep track of the data encryption keys used to encrypt data, as they are stored with the encrypted data and can be retrieved and decrypted by using AWS KMS when needed. Option A is not optimal because it will require manual tracking of the data encryption keys used for each data object, which is error-prone and inefficient. Option C is not optimal because it will store the data encryption keys automatically in Amazon S3, which is unnecessary and insecure as Amazon S3 is not designed for storing encryption keys. Option D is not optimal because it will store the data encryption key in the user data for the EC2 instance, which is also unnecessary and insecure as user data is not encrypted by default.
Reference: [AWS Encryption SDK], [AWS Key Management Service]

QUESTION 3
A company stores customer credit reports in an Amazon S3 bucket. An analytics service uses standard Amazon S3 GET requests to access the reports. A developer must implement a solution to redact personally identifiable information (PII) from the reports before the reports reach the analytics service.

A. Load the S3 objects into Amazon Redshift by using a COPY command. Implement dynamic data masking. Refactor the analytics service to read from Amazon Redshift.
B. Set up an S3 Object Lambda function. Attach the function to an S3 Object Lambda Access Point. Program the function to call a PII redaction API.
C. Use AWS Key Management Service (AWS KMS) to implement encryption in the S3 bucket. Re-upload all the existing S3 objects. Give the kms permission to the analytics service.
D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Implement message data protection. Refactor the analytics service to publish data access requests to the SNS topic.

Correct Answer: B
Section:
Explanation:
Comprehensive Detailed Step by Step Explanation with All AWS Developer
Reference:
To redact PII from S3 objects before they are accessed by the analytics service, the most efficient solution is to use S3 Object Lambda. S3 Object Lambda allows you to add your own code (Lambda function) to process and transform data when it is retrieved from Amazon S3. You can attach a Lambda function to an S3 Object Lambda Access Point, which in this case would run a redaction API to remove PII from the reports.
Operational Efficiency: S3 Object Lambda handles data processing on the fly, without requiring the data to be permanently transformed or moved to another service (like Amazon Redshift).
Alternatives:
Option A: Loading the data into Amazon Redshift would require refactoring the analytics service and maintaining an additional data pipeline, increasing complexity.
Option C: Using AWS KMS for encryption protects data at rest and in transit, but it does not address PII redaction.
Option D: SNS is a messaging service and does not support direct data transformation.

QUESTION 4
A company is using the AWS Serverless Application Model (AWS SAM) to develop a social media application. A developer needs a quick way to test AWS Lambda functions locally by using test event payloads. The developer needs the structure of these test event payloads to match the actual events that AWS services create.

A. Create shareable test Lambda events. Use these test Lambda events for local testing.
B. Store manually created test event payloads locally. Use the sam local invoke command with the file path to the payloads.
C. Store manually created test event payloads in an Amazon S3 bucket. Use the sam local invoke command with the S3 path to the payloads.
D. Use the sam local generate-event command to create test payloads for local testing.

Correct Answer: D
Section:
Explanation:
Comprehensive Detailed Step by Step Explanation with All AWS Developer
Reference:
The AWS Serverless Application Model (SAM) includes features for local testing and debugging of AWS Lambda functions. One of the most efficient ways to generate test payloads that match actual AWS event structures is by using the sam local generate-event command.
sam local generate-event: This command allows developers to create pre-configured test event payloads for various AWS services (e.g., S3, API Gateway, SNS). These generated events accurately reflect the format that the service would use in a live environment, reducing the manual work required to create these events from scratch.
Operational Overhead: This approach reduces overhead since the developer does not need to manually create or maintain test events. It ensures that the structure is correct and up-to-date with the latest AWS standards.
Alternatives:
Option A suggests using shareable test events, but manually creating or sharing these events introduces more overhead.
Option B and C both involve manually storing and maintaining test events, which adds unnecessary complexity compared to using sam local generate-event.
AWS SAM CLI documentation

QUESTION 5
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 me S3 bucket in the MOST secure way?

A. Hardcode the credentials that are required to access the S3 objects in the application code. Use the credentials to access me required S3 objects.

B. Create a secret access key and access key ID with permission to access the S3 bucket. Store the key and key ID in AWS Secrets Manager. Configure the application to retrieve the Secrets Manager secret and use the credentials to access me S3 objects.
C. Create a Lambda function execution role Attach a policy to the rote that grants access to specific objects in the S3 bucket.
D. 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 Lambda. Use the environment variables to access the required S3 objects.

Correct Answer: C
Section:
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.
Reference: [AWS Lambda Execution Role], [Using AWS Lambda with Amazon S3]

QUESTION 6
A developer has code that is stored in an Amazon S3 bucket. The code must be deployed as an AWS Lambda function across multiple accounts in the same AWS Region as the S3 bucket an AWS CloudPormation template that runs for each account will deploy the Lambda function.
What is the MOST secure way to allow CloudFormaton to access the Lambda Code in the S3 bucket?

A. Grant the CloudFormation service role the S3 ListBucket and GetObject permissions. Add a bucket policy to Amazon S3 with the principal of "AWS" (account numbers)
B. Grant the CloudFormation service row the S3 GetObfect permission. Add a Bucket policy to Amazon S3 with the principal of """
C. Use a service-based link to grant the Lambda function the S3 ListBucket and GetObject permissions by explicitly adding the S3 bucket's account number in the resource.
D. Use a service-based link to grant the Lambda function the S3 GetObject permission Add a resource of "** to allow access to the S3 bucket.

Correct Answer: B
Section:
Explanation:
This solution allows the CloudFormation service role to access the S3 bucket from any account, as long as it has the S3 GetObject permission. The bucket policy grants access to any principal with the GetObject permission, which is the least privilege needed to deploy the Lambda code. This is more secure than granting ListBucket permission, which is not required for deploying Lambda code, or using a service-based link, which is not supported for Lambda functions.
Reference: AWS CloudFormation Service Role, Using AWS Lambda with Amazon S3

QUESTION 7
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 file. Create a new API Import the OpenAPI file Modify the new API to add request validation. Perform the tests Modify the existing API to add request validation. Deploy the existing API to production.
B. Modify the existing API to add request validation. Deploy the updated API to a new API Gateway stage Perform the tests Deploy the updated API to the API Gateway production stage.
C. Create a new API Add the necessary resources and methods including new request validation. Perform the tests Modify the existing API to add request validation. Deploy the existing API to production.
D. Clone the exiting API Modify the new API lo add request validation. Perform the tests Modify the existing API to add request validation Deploy the existing API to production.

Correct Answer: D
Section:
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]

QUESTION 8
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 account. Add an EventBridge rule to the event bus of the main account that matches all EC2 instance lifecycle events. Add the SQS queue as a target of the rule.
B. Use the resource policies of the SQS queue in the main account to give each account permissions to write to that SQS queue. Add to the Amazon EventBridge event bus of each account an EventBridge rule that matches all EC2 instance lifecycle events. Add the SQS queue in the main account as a target of the rule.
C. Write an AWS Lambda function that scans through all EC2 instances in the company accounts to detect EC2 instance lifecycle changes. 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 change. Add an Amazon EventBridge scheduled rule that invokes the Lambda function every minute.
D. Configure the permissions on the main account event bus to receive events from all accounts. Create an Amazon EventBridge rule in each account to send all the EC2 instance lifecycle events to the main account event bus. Add an EventBridge rule to the main account event bus that matches all EC2 instance lifecycle events. Set the SQS queue as a target for the rule.

Correct Answer: D
Section:
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

QUESTION 9
An application is using Amazon Cognito user pools and identity pools for secure access. A developer wants to integrate the user-specific file upload and download features in the application with Amazon S3. The developer must ensure that the files are saved and retrieved in a secure manner and that users can access only their own files. The file sizes range from 3 KB to 300 MB.
Which option will meet these requirements with the HIGHEST level of security?

A. Use S3 Event Notifications to validate the file upload and download requests and update the user interface (UI).
B. Save the details of the uploaded files in a separate Amazon DynamoDB table. Filter the list of files in the user interface (UI) by comparing the current user ID with the user ID associated with the file in the table.
C. Use Amazon API Gateway and an AWS Lambda function to upload and download files. Validate each request in the Lambda function before performing the requested operation.
D. Use an IAM policy within the Amazon Cognito identity prefix to restrict users to use their own folders in Amazon S3.

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/cognito/latest/developerguide/amazon-cognito-integrating-userpools-with-identity-pools.html

QUESTION 10
A company is building a scalable data management solution by using AWS services to improve the speed and agility of development. The solution will ingest large volumes of data from various sources and will process this data through multiple business rules and transformations.
The solution requires business rules to run in sequence and to handle reprocessing of data if errors occur when the business rules run. The company needs the solution to be scalable and to require the least possible maintenance.
Which AWS service should the company use to manage and automate the orchestration of the data flows to meet these requirements?

A. AWS Batch
B. AWS Step Functions

C. AWS Glue
D. AWS Lambda

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/step-functions/latest/dg/welcome.html

QUESTION 11
A developer has created an AWS Lambda function that is written in Python. The Lambda function reads data from objects in Amazon S3 and writes data to an Amazon DynamoDB table. The function is successfully invoked from an S3 event notification when an object is created. However, the function fails when it attempts to write to the DynamoDB table.
What is the MOST likely cause of this issue?

A. The Lambda function's concurrency limit has been exceeded.
B. DynamoDB table requires a global secondary index (GSI) to support writes.
C. The Lambda function does not have IAM permissions to write to DynamoDB.
D. The DynamoDB table is not running in the same Availability Zone as the Lambda function.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_lambda-access-dynamodb.html

QUESTION 12
Users are reporting errors in an application. The application consists of several micro services that are deployed on Amazon Elastic Container Serves (Amazon ECS) with AWS Fargate.
When combination of steps should a developer take to fix the errors? (Select TWO)

A. Deploy AWS X-Ray as a sidecar container to the micro services. Update the task role policy to allow access to me X -Ray API.
B. Deploy AWS X-Ray as a daemon set to the Fargate cluster. Update the service role policy to allow access to the X-Ray API.
C. Instrument the application by using the AWS X-Ray SDK. Update the application to use the Put-XrayTrace API call to communicate with the X-Ray API.
D. Instrument the application by using the AWS X-Ray SDK. Update the application to communicate with the X-Ray daemon.
E. Instrument the ECS task to send the stout and spider- output to Amazon CloudWatch Logs. Update the task role policy to allow the cloudwatch Putlogs action.

Correct Answer: A, E
Section:
Explanation:
The combination of steps that the developer should take to fix the errors is to deploy AWS X-Ray as a sidecar container to the microservices and instrument the ECS task to send the stdout and stderr output to Amazon CloudWatch Logs. This way, the developer can use AWS X-Ray to analyze and debug the performance of the microservices and identify any issues or bottlenecks. The developer can also use CloudWatch Logs to monitor and troubleshoot the logs from the ECS task and detect any errors or exceptions. The other options either involve using AWS X-Ray as a daemon set, which is not supported by Fargate, or using the PutTraceSegments API call, which is not necessary when using a sidecar container.
Reference: Using AWS X-Ray with Amazon ECS

QUESTION 13
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.

Correct Answer: C
Section:
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]

QUESTION 14
A developer is building a serverless application that is based on AWS Lambd a. The developer initializes the AWS software development kit (SDK) outside of the Lambda handcar function.
What is the PRIMARY benefit of this action?

A. Improves legibility and systolic convention
B. Takes advantage of runtime environment reuse
C. Provides better error handling
D. Creates a new SDK instance for each invocation

Correct Answer: B
Section:
Explanation:
This benefit occurs when initializing the AWS SDK outside of the Lambda handler function because it allows the SDK instance to be reused across multiple invocations of the same function. This can improve performance and reduce latency by avoiding unnecessary initialization overhead. If the SDK is initialized inside the handler function, it will create a new SDK instance for each invocation, which can increase memory usage and execution time.
Reference: [AWS Lambda execution environment], [Best Practices for Working with AWS Lambda Functions]

QUESTION 15
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)

Correct Answer: C
Section:
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]

QUESTION 16
A developer at a company recently created a serverless application to process and show data from business reports. The application's user interface (UI) allows users to select and start processing the files. The Ul displays a message when the result is available to view. The application uses AWS Step Functions with AWS Lambda functions to process the files. The developer used Amazon API Gateway and Lambda functions to create an API to

support the UI.
The company's Ul team reports that the request to process a file is often returning timeout errors because of the see or complexity of the files. The Ul team wants the API to provide an immediate response so that the Ul can deploy a message while the files are being processed. The backend process that is invoked by the API needs to send an email message when the report processing is complete.
What should the developer do to configure the API to meet these requirements?

A. Change the API Gateway route to add an X-Amz-Invocation-Type header win a sialic value of 'Event' in the integration request Deploy the API Gateway stage to apply the changes.
B. Change the configuration of the Lambda function that implements the request to process a file.
Configure the maximum age of the event so that the Lambda function will ion asynchronously.
C. Change the API Gateway timeout value to match the Lambda function ominous value. Deploy the API Gateway stage to apply the changes.
D. Change the API Gateway route to add an X-Amz-Target header with a static value of 'A sync' in the integration request Deploy me API Gateway stage to apply the changes.

Correct Answer: A
Section:
Explanation:
This solution allows the API to invoke the Lambda function asynchronously, which means that the API will return an immediate response without waiting for the function to complete. The X-Amz-Invocation-Type header specifies the invocation type of the Lambda function, and setting it to 'Event' means that the function will be invoked asynchronously. The function can then use Amazon Simple Email Service (SES) to send an email message when the report processing is complete.
Reference: [Asynchronous invocation], [Set up Lambda proxy integrations in API Gateway]

QUESTION 17
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.

Correct Answer: C
Section:
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]

QUESTION 18
A mobile app stores blog posts in an Amazon DynacnoDB table Millions of posts are added every day and each post represents a single item in the table. The mobile app requires only recent posts. Any post that is older than 48 hours can be removed.
What is the MOST cost-effective way to delete posts that are older man 48 hours?

A. For each item add a new attribute of type String that has a timestamp that is set to the blog post creation time. Create a script to find old posts with a table scan and remove posts that are order than 48 hours by using the Balch Write Item API operation. Schedule a cron job on an Amazon EC2 instance once an hour to start the script.
B. For each item add a new attribute of type. String that has a timestamp that its set to the blog post creation time. Create a script to find old posts with a table scan and remove posts that are Oder than 48 hours by using the Batch Write item API operating. Place the script in a container image. Schedule an Amazon Elastic Container Service (Amazon ECS) task on AWS Far gate that invokes the container every 5 minutes.
C. For each item, add a new attribute of type Date that has a timestamp that is set to 48 hours after the blog post creation time. Create a global secondary index (GSI) that uses the new attribute as a sort key. Create an AWS Lambda function that references the GSI and removes expired items by using the Batch Write item API operation Schedule me function with an Amazon CloudWatch event every minute.

D. For each item add a new attribute of type. Number that has timestamp that is set to 48 hours after the blog post. creation time Configure the DynamoDB table with a TTL that references the new attribute.

Correct Answer: D
Section:
Explanation:
This solution will meet the requirements by using the Time to Live (TTL) feature of DynamoDB, which enables automatically deleting items from a table after a certain time period. The developer can add a new attribute of type Number that has a timestamp that is set to 48 hours after the blog post creation time, which represents the expiration time of the item. The developer can configure the DynamoDB table with a TTL that references the new attribute, which instructs DynamoDB to delete the item when the current time is greater than or equal to the expiration time. This solution is also cost-effective as it does not incur any additional charges for deleting expired items. Option A is not optimal because it will create a script to find and remove old posts with a table scan and a batch write item API operation, which may consume more read and write capacity units and incur more costs. Option B is not optimal because it will use Amazon Elastic Container Service (Amazon ECS) and AWS Fargate to run the script, which may introduce additional costs and complexity for managing and scaling containers. Option C is not optimal because it will create a global secondary index (GSI) that uses the expiration time as a sort key, which may consume more storage space and incur more costs.
Reference: Time To Live, Managing DynamoDB Time To Live (TTL)

QUESTION 19
A developer is modifying an existing AWS Lambda function White checking the code the developer notices hardcoded parameter various for an Amazon RDS for SQL Server user name password database host and port. There also are hardcoded parameter values for an Amazon DynamoOB table.
an Amazon S3 bucket, and an Amazon Simple Notification Service (Amazon SNS) topic.
The developer wants to securely store the parameter values outside the code m an encrypted format and wants to turn on rotation for the credentials. The developer also wants to be able to reuse the parameter values from other applications and to update the parameter values without modifying code.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an RDS database secret in AWS Secrets Manager. Set the user name password, database, host and port. Turn on secret rotation. Create encrypted Lambda environment variables for the DynamoDB table, S3 bucket and SNS topic.
B. Create an RDS database secret in AWS Secrets Manager. Set the user name password, database, host and port. Turn on secret rotation. Create Secure String parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket and SNS topic.
C. Create RDS database parameters in AWS Systems Manager Parameter. Store for the user name password, database, host and port. Create encrypted Lambda environment variables for me DynamoDB table, S3 bucket, and SNS topic. Create a Lambda function and set the logic for the credentials rotation task Schedule the credentials rotation task in Amazon EventBridge.
D. Create RDS database parameters in AWS Systems Manager Parameter. Store for the user name password database, host, and port. Store the DynamoDB table. S3 bucket, and SNS topic in Amazon S3 Create a Lambda function and set the logic for the credentials rotation Invoke the Lambda function on a schedule.

Correct Answer: B
Section:
Explanation:
This solution will meet the requirements by using AWS Secrets Manager and AWS Systems Manager Parameter Store to securely store the parameter values outside the code in an encrypted format. AWS Secrets Manager is a service that helps protect secrets such as database credentials by encrypting them with AWS Key Management Service (AWS KMS) and enabling automatic rotation of secrets. The developer can create an RDS database secret in AWS Secrets Manager and set the user name, password, database, host, and port for accessing the RDS database. The developer can also turn on secret rotation, which will change the database credentials periodically according to a specified schedule or event. AWS Systems Manager Parameter Store is a service that provides secure and scalable storage for configuration data and secrets. The developer can create Secure String parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket, and SNS topic, which will encrypt them with AWS KMS. The developer can also reuse the parameter values from other applications and update them without modifying code. Option A is not optimal because it will create encrypted Lambda environment variables for the DynamoDB table, S3 bucket, and SNS topic, which may not be reusable or updatable without modifying code. Option C is not optimal because it will create RDS database parameters in AWS Systems Manager Parameter Store, which does not support automatic rotation of secrets. Option D is not optimal because it will store the DynamoDB table, S3 bucket, and SNS topic in Amazon S3, which may introduce additional costs and complexity for accessing configuration data.
Reference: AWS Secrets Manager, [AWS Systems Manager Parameter Store]

QUESTION 20
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.

Correct Answer: D
Section:
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.

QUESTION 21
A developer has an application that makes batch requests directly to Amazon DynamoDB by using the BatchGetItem low-level API operation. The responses frequently return values in the UnprocessedKeys element.
Which actions should the developer take to increase the resiliency of the application when the batch response includes values in UnprocessedKeys? (Choose two.)

A. Retry the batch operation immediately.
B. Retry the batch operation with exponential backoff and randomized delay.
C. Update the application to use an AWS software development kit (AWS SDK) to make the requests.
D. Increase the provisioned read capacity of the DynamoDB tables that the operation accesses.
E. Increase the provisioned write capacity of the DynamoDB tables that the operation accesses.

Correct Answer: B, C
Section:
Explanation:
The UnprocessedKeys element indicates that the BatchGetItem operation did not process all of the requested items in the current response. This can happen if the response size limit is exceeded or if the table's provisioned throughput is exceeded. To handle this situation, the developer should retry the batch operation with exponential backoff and randomized delay to avoid throttling errors and reduce the load on the table. The developer should also use an AWS SDK to make the requests, as the SDKs automatically retry requests that return UnprocessedKeys.
Reference:
[BatchGetItem - Amazon DynamoDB]
[Working with Queries and Scans - Amazon DynamoDB]
[Best Practices for Handling DynamoDB Throttling Errors]

QUESTION 22
A company is running a custom application on a set of on-premises Linux servers that are accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled on the API test stage.
How can a developer enable X-Ray tracing on the on-premises servers with the LEAST amount of configuration?

A. Install and run the X-Ray SDK on the on-premises servers to capture and relay the data to the XRay service.
B. Install and run the X-Ray daemon on the on-premises servers to capture and relay the data to the X-Ray service.
C. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTraceSegments API call.
D. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTelemetryRecords API call.

Correct Answer: B
Section:
Explanation:
The X-Ray daemon is a software that collects trace data from the X-Ray SDK and relays it to the X-Ray service. The X-Ray daemon can run on any platform that supports Go, including Linux, Windows, and macOS. The developer can install and run the X-Ray daemon on the on-premises servers to capture and relay the data to the X-Ray service with minimal configuration. The X-Ray SDK is used to instrument the application code, not to capture and relay data. The Lambda function solutions are more complex and require additional configuration.
Reference:
[AWS X-Ray concepts - AWS X-Ray]
[Setting up AWS X-Ray - AWS X-Ray]

QUESTION 23
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. Store the API credentials in AWS Secrets Manager. Retrieve the API credentials at runtime by using the AWS SDK. Use the credentials to make the API call.
B. Store the API credentials in a local code variable. Push the code to a secure Git repository. Use the local code variable at runtime to make the API call.
C. Store the API credentials as an object in a private Amazon S3 bucket. Restrict access to the S3 object by using IAM policies. Retrieve the API credentials at runtime by using the AWS SDK. Use the credentials to make the API call.
D. Store the API credentials in an Amazon DynamoDB table. Restrict access to the table by using resource-based policies. Retrieve the API credentials at runtime by using the AWS SDK. Use the credentials to make the API call.

Correct Answer: A
Section:
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.
Reference:
[What Is AWS Secrets Manager? - AWS Secrets Manager]
[Retrieving a Secret - AWS Secrets Manager]

QUESTION 24
A developer is deploying a new application to Amazon Elastic Container Service (Amazon ECS). The developer needs to securely store and retrieve different types of variables. These variables include authentication information for a remote API, the URL for the API, and credentials. The authentication information and API URL must be available to all current and future deployed versions of the application across development, testing, and production environments.
How should the developer retrieve the variables with the FEWEST application changes?

A. Update the application to retrieve the variables from AWS Systems Manager Parameter Store. Use unique paths in Parameter Store for each variable in each environment. Store the credentials in AWS Secrets Manager in each environment.
B. Update the application to retrieve the variables from AWS Key Management Service (AWS KMS).
Store the API URL and credentials as unique keys for each environment.
C. Update the application to retrieve the variables from an encrypted file that is stored with the application. Store the API URL and credentials in unique files for each environment.
D. Update the application to retrieve the variables from each of the deployed environments. Define the authentication information and API URL in the ECS task definition as unique names during the deployment process.

Correct Answer: A
Section:
Explanation:
AWS Systems Manager Parameter Store is a service that provides secure, hierarchical storage for configuration data management and secrets management. The developer can update the application to retrieve the variables from Parameter Store by using the AWS SDK or the AWS CLI. The developer can use unique paths in Parameter Store for each variable in each environment, such as /dev/api-url, /test/api-url, and /prod/api-url. The developer can also store the credentials in AWS Secrets Manager, which is integrated with Parameter Store and provides additional features such as automatic rotation and encryption.
Reference:
[What Is AWS Systems Manager? - AWS Systems Manager]
[Parameter Store - AWS Systems Manager]
[What Is AWS Secrets Manager? - AWS Secrets Manager]

QUESTION 25
A company is migrating legacy internal applications to AWS. Leadership wants to rewrite the internal employee directory to use native AWS services. A developer needs to create a solution for storing employee contact details

and high-resolution photos for use with the new application.
Which solution will enable the search and retrieval of each employee's individual details and highresolution photos using AWS APIs?

A. Encode each employee's contact information and photos using Base64. Store the information in an Amazon DynamoDB table using a sort key.
B. Store each employee's contact information in an Amazon DynamoDB table along with the object keys for the photos stored in Amazon S3.
C. Use Amazon Cognito user pools to implement the employee directory in a fully managed software-as-a-service (SaaS) method.
D. Store employee contact information in an Amazon RDS DB instance with the photos stored in Amazon Elastic File System (Amazon EFS).

Correct Answer: B
Section:
Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance with seamless scalability. The developer can store each employee's contact information in a DynamoDB table along with the object keys for the photos stored in Amazon S3. Amazon S3 is an object storage service that offers industry-leading scalability, data availability, security, and performance. The developer can use AWS APIs to search and retrieve the employee details and photos from DynamoDB and S3.
Reference:
[Amazon DynamoDB]
[Amazon Simple Storage Service (S3)]

QUESTION 26
A company is expanding the compatibility of its photo-snaring mobile app to hundreds of additional devices with unique screen dimensions and resolutions. Photos are stored in Amazon S3 in their original format and resolution. The company uses an Amazon CloudFront distribution to serve the photos The app includes the dimension and resolution of the display as GET parameters with every request.
A developer needs to implement a solution that optimizes the photos that are served to each device to reduce load time and increase photo quality.
Which solution will meet these requirements MOST cost-effective?

A. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolutions. Create a dynamic CloudFront origin that automatically maps the request of each device to the corresponding photo variant.
B. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolutions. Create a Lambda@Edge function to route requests to the corresponding photo vacant by using request headers.
C. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response. Change the CloudFront TTL cache policy to the maximum value possible.
D. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response. In the same function store a copy of the processed photos on Amazon S3 for subsequent requests.

Correct Answer: D
Section:
Explanation:
This solution meets the requirements most cost-effectively because it optimizes the photos on demand and caches them for future requests. Lambda@Edge allows the developer to run Lambda functions at AWS locations closer to viewers, which can reduce latency and improve photo quality. The developer can create a Lambda@Edge function that uses the GET parameters from each request to optimize the photos with the required dimensions and resolutions and returns them as a response. The function can also store a copy of the processed photos on Amazon S3 for subsequent requests, which can reduce processing time and costs. Using S3 Batch Operations to create new variants of the photos will incur additional storage costs and may not cover all possible dimensions and resolutions. Creating a dynamic CloudFront origin or a Lambda@Edge function to route requests to corresponding photo variants will require maintaining a mapping of device types and photo variants, which can be complex and error-prone.
Reference: [Lambda@Edge Overview], [Resizing Images with Amazon CloudFront & Lambda@Edge]

QUESTION 27
A developer has been asked to create an AWS Lambda function that is invoked any time updates are made to items in an Amazon DynamoDB table. The function has been created and appropriate permissions have been added to the Lambda execution role Amazon DynamoDB streams have been enabled for the table, but the function 15 still not being invoked.
Which option would enable DynamoDB table updates to invoke the Lambda function?

A. Change the StreamViewType parameter value to NEW_AND_OLOJMAGES for the DynamoDB table.
B. Configure event source mapping for the Lambda function.

C. Map an Amazon Simple Notification Service (Amazon SNS) topic to the DynamoDB streams.
D. Increase the maximum runtime (timeout) setting of the Lambda function.

Correct Answer: B
Section:
Explanation:
This solution allows the Lambda function to be invoked by the DynamoDB stream whenever updates are made to items in the DynamoDB table. Event source mapping is a feature of Lambda that enables a function to be triggered by an event source, such as a DynamoDB stream, an Amazon Kinesis stream, or an Amazon Simple Queue Service (SQS) queue. The developer can configure event source mapping for the Lambda function using the AWS Management Console, the AWS CLI, or the AWS SDKs. Changing the StreamViewType parameter value to NEW_AND_OLD_IMAGES for the DynamoDB table will not affect the invocation of the Lambda function, but only change the information that is written to the stream record. Mapping an Amazon Simple Notification Service (Amazon SNS) topic to the DynamoDB stream will not invoke the Lambda function directly, but require an additional subscription from the Lambda function to the SNS topic. Increasing the maximum runtime (timeout) setting of the Lambda function will not affect the invocation of the Lambda function, but only change how long the function can run before it is terminated.
Reference: [Using AWS Lambda with Amazon DynamoDB], [Using AWS Lambda with Amazon SNS]

QUESTION 28
A company is using Amazon OpenSearch Service to implement an audit monitoring system. A developer needs to create an AWS Cloudformation custom resource that is associated with an AWS Lambda function to configure the OpenSearch Service domain. The Lambda function must access the OpenSearch Service domain by using Open Search Service internal master user credentials.
What is the MOST secure way to pass these credentials to the Lambdas function?

A. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment variable. Set the No Echo attenuate to true.
B. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and to create a parameter. In AWS Systems Manager Parameter Store. Set the No Echo attribute to true. Create an 1AM role that has the ssm GetParameter permission. Assign me role to the Lambda function. Store me parameter name as the Lambda function's environment variable. Resolve the parameter's value at runtime.
C. Use a CloudFormation parameter to pass the master uses credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment varleWe Encrypt the parameters value by using the AWS Key Management Service (AWS KMS) encrypt command.
D. Use CloudFoimalion to create an AWS Secrets Manager Secret. Use a CloudFormation dynamic reference to retrieve the secret's value for the OpenSearch Service domain's MasterUserOptions. Create an 1AM role that has the secrets manager. GetSecretvalue permission. Assign the role to the Lambda Function Store the secrets name as the Lambda function's environment variable. Resole the secret's value at runtime.

Correct Answer: D
Section:
Explanation:
The solution that will meet the requirements is to use CloudFormation to create an AWS Secrets Manager secret. Use a CloudFormation dynamic reference to retrieve the secret's value for the OpenSearch Service domain's MasterUserOptions. Create an IAM role that has the secretsmanager:GetSecretValue permission. Assign the role to the Lambda function. Store the secret's name as the Lambda function's environment variable. Resolve the secret's value at runtime.
This way, the developer can pass the credentials to the Lambda function in a secure way, as AWS Secrets Manager encrypts and manages the secrets. The developer can also use a dynamic reference to avoid exposing the secret's value in plain text in the CloudFormation template. The other options either involve passing the credentials as plain text parameters, which is not secure, or encrypting them with AWS KMS, which is less convenient than using AWS Secrets Manager.
Reference: Using dynamic references to specify template values

QUESTION 29
An application runs on multiple EC2 instances behind an ELB.
Where is the session data best written so that it can be served reliably across multiple requests?

A. Write data to Amazon ElastiCache
B. Write data to Amazon Elastic Block Store
C. Write data to Amazon EC2 instance Store
D. Wide data to the root filesystem

Correct Answer: A
Section:
Explanation:
The solution that will meet the requirements is to write data to Amazon ElastiCache. This way, the application can write session data to a fast, scalable, and reliable in-memory data store that can be served reliably across multiple requests. The other options either involve writing data to persistent storage, which is slower and more expensive than in-memory storage, or writing data to the root filesystem, which is not shared among multiple EC2 instances.
Reference: Using ElastiCache for session management

QUESTION 30
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

Correct Answer: C
Section:
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

QUESTION 31
A developer migrated a legacy application to an AWS Lambda function. The function uses a thirdparty 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 failures. Use the Wait state to delay calling the Lambda function.
B. Use an Amazon Simple Queue Service (Amazon SQS) queue to hold the API calls. Configure the Lambda function to poll the queue within the API threshold limits.
C. Use an Amazon CloudWatch Logs metric to count the number of API calls. Configure an Amazon CloudWatch alarm flat slops the currently running instance of the Lambda function when the metric exceeds the API threshold limits.
D. 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.

Correct Answer: A
Section:
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

QUESTION 32
A developer must analyze performance issues with production-distributed applications written as AWS Lambda functions. These distributed Lambda applications invoke other components that make up me applications. How

should the developer identify and troubleshoot the root cause of the performance issues in production?

A. Add logging statements to the Lambda functions. then use Amazon CloudWatch to view the logs.
B. Use AWS CloudTrail and then examine the logs.
C. Use AWS X-Ray. then examine the segments and errors.
D. Run Amazon inspector agents and then analyze performance.

Correct Answer: C
Section:
Explanation:
This solution will meet the requirements by using AWS X-Ray to analyze and debug the performance issues with the distributed Lambda applications. 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 use AWS X-Ray to identify the root cause of the performance issues by examining the segments and errors that show the details of each request and the components that make up the applications. Option A is not optimal because it will use logging statements and Amazon CloudWatch, which may not provide enough information or visibility into the distributed applications. Option B is not optimal because it will use AWS CloudTrail, which is a service that records API calls and events for AWS services, not application performance data. Option D is not optimal because it will use Amazon Inspector, which is a service that helps improve the security and compliance of applications on Amazon EC2 instances, not Lambda functions.
Reference: AWS X-Ray, Using AWS X-Ray with AWS Lambda

QUESTION 33
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. Bluegreen
D. Immutable

Correct Answer: B
Section:
Explanation:
This solution will meet the requirements by using a rolling with additional batch deployment method, which deploys the new version of the application to a separate group of instances and then shifts traffic to those instances in batches. This way, the application maintains full capacity and avoids service interruption during deployment, as well as minimizes the cost of additional resources that support the deployment. Option A is not optimal because it will use an all at once deployment method, which deploys the new version of the application to all instances simultaneously, which may cause service interruption or downtime during deployment. Option C is not optimal because it will use a blue/green deployment method, which deploys the new version of the application to a separate environment and then swaps URLs with the original environment, which may incur more costs for additional resources that support the deployment. Option D is not optimal because it will use an immutable deployment method, which deploys the new version of the application to a fresh group of instances and then redirects traffic to those instances, which may also incur more costs for additional resources that support the deployment.
Reference: AWS Elastic Beanstalk Deployment Policies

QUESTION 34
A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node is application. To minimize these bugs, the developer wants to impendent automated testing of Lambda functions in an environment that Closely simulates the Lambda environment.
The developer needs to give other developers the ability to run the tests locally. The developer also needs to integrate the tests into the team's continuous integration and continuous delivery (Ct/CO) pipeline before the AWS Cloud Development Kit (AWS COK) deployment.
Which solution will meet these requirements?

A. Create sample events based on the Lambda documentation. Create automated test scripts that use the cdk local invoke command to invoke the Lambda functions. Check the response Document the test scripts for the other developers on the team Update the CI/CD pipeline to run the test scripts.
B. Install a unit testing framework that reproduces the Lambda execution environment. Create sample events based on the Lambda Documentation Invoke the handler function by using a unit testing framework. Check the

response Document how to run the unit testing framework for the other developers on the team. Update the OCD pipeline to run the unit testing framework.

C. Install the AWS Serverless Application Model (AWS SAW) CLI tool Use the Sam local generateevent command to generate sample events for me automated tests. Create automated test scripts that use the Sam local invoke command to invoke the Lambda functions. Check the response
Document the test scripts tor the other developers on the team Update the CI/CD pipeline to run the test scripts.

D. Create sample events based on the Lambda documentation. Create a Docker container from the Node is base image to invoke the Lambda functions. Check the response Document how to run the Docker container for the more developers on the team update the CI/CD pipeline to run the Docker container.

Correct Answer: C
Section:
Explanation:
This solution will meet the requirements by using AWS SAM CLI tool, which is a command line tool that lets developers locally build, test, debug, and deploy serverless applications defined by AWS SAM templates. The developer can use sam local generate-event command to generate sample events for different event sources such as API Gateway or S3. The developer can create automated test scripts that use sam local invoke command to invoke Lambda functions locally in an environment that closely simulates Lambda environment. The developer can check the response from Lambda functions and document how to run the test scripts for other developers on the team. The developer can also update CI/CD pipeline to run these test scripts before deploying with AWS CDK. Option A is not optimal because it will use cdk local invoke command, which does not exist in AWS CDK CLI tool.
Option B is not optimal because it will use a unit testing framework that reproduces Lambda execution environment, which may not be accurate or consistent with Lambda environment. Option D is not optimal because it will create a Docker container from Node.js base image to invoke Lambda functions, which may introduce additional overhead and complexity for creating and running Docker containers.
Reference: [AWS Serverless Application Model (AWS SAM)], [AWS Cloud Development Kit (AWS CDK)]

QUESTION 35
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 principal. Associate the role with the policy.
D. Create a trust relationship between the role and dynamodb Amazonas com.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/access-controloverview.html#access-control-resource-ownership

QUESTION 36
A developer needs to deploy an application running on AWS Fargate using Amazon ECS The application has environment variables that must be passed to a container for the application to initialize.

How should the environment variables be passed to the container?

A. Define an array that includes the environment variables under the environment parameter within the service definition.
B. Define an array that includes the environment variables under the environment parameter within the task definition.
C. Define an array that includes the environment variables under the entryPoint parameter within the task definition.
D. Define an array that includes the environment variables under the entryPoint parameter within the service definition.

Correct Answer: B
Section:
Explanation:
This solution allows the environment variables to be passed to the container when it is launched by AWS Fargate using Amazon ECS. The task definition is a text file that describes one or more containers that form an application. It contains various parameters for configuring the containers, such as CPU and memory requirements, network mode, and environment variables. The environment parameter is an array of key-value pairs that specify environment variables to pass to a container. Defining an array that includes the environment variables under the entryPoint parameter within the task definition will not pass them to the container, but use them as command-line arguments for overriding the default entry point of a container. Defining an array that includes the environment variables under the environment or entryPoint parameter within the service definition will not pass them to the container, but cause an error because these parameters are not valid for a service definition.
Reference: [Task Definition Parameters], [Environment Variables]

QUESTION 37
A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the AWS Key Management Service (AWS KMS) key was used and by whom.
Which encryption option will meet these requirements?

A. Server-side encryption with Amazon S3 managed keys (SSE-S3)
B. Server-side encryption with AWS KMS managed keys (SSE-KMS}
C. Server-side encryption with customer-provided keys (SSE-C)
D. Server-side encryption with self-managed keys

Correct Answer: B
Section:
Explanation:
This solution meets the requirements because it encrypts data at rest using AWS KMS keys and provides an audit trail of when and by whom they were used. Server-side encryption with AWS KMS managed keys (SSE-KMS) is a feature of Amazon S3 that encrypts data using keys that are managed by AWS KMS. When SSE-KMS is enabled for an S3 bucket or object, S3 requests AWS KMS to generate data keys and encrypts data using these keys. AWS KMS logs every use of its keys in AWS CloudTrail, which records all API calls to AWS KMS as events. These events include information such as who made the request, when it was made, and which key was used. The company policy can use CloudTrail logs to audit critical events related to their data encryption and access. Server-side encryption with Amazon S3 managed keys (SSE-S3) also encrypts data at rest using keys that are managed by S3, but does not provide an audit trail of key usage. Server-side encryption with customer-provided keys (SSE-C) and server-side encryption with self-managed keys also encrypt data at rest using keys that are provided or managed by customers, but do not provide an audit trail of key usage and require additional overhead for key management.
Reference: [Protecting Data Using Server-Side Encryption with AWS KMS-Managed Encryption Keys (SSE-KMS)], [Logging AWS KMS API calls with AWS CloudTrail]

QUESTION 38
A company is building a new application that runs on AWS and uses Amazon API Gateway to expose APIs Teams of developers are working on separate components of the application in parallel The company wants to publish an API without an integrated backend so that teams that depend on the application backend can continue the development work before the API backend development is complete.
Which solution will meet these requirements?

A. Create API Gateway resources and set the integration type value to MOCK Configure the method integration request and integration response to associate a response with an HTTP status code Create an API Gateway stage and deploy the API.
B. Create an AWS Lambda function that returns mocked responses and various HTTP status codes. Create API Gateway resources and set the integration type value to AWS_PROXY Deploy the API.
C. Create an EC2 application that returns mocked HTTP responses Create API Gateway resources and set the integration type value to AWS Create an API Gateway stage and deploy the API.
D. Create API Gateway resources and set the integration type value set to HTTP_PROXY. Add mapping templates and deploy the API. Create an AWS Lambda layer that returns various HTTP status codes Associate the Lambda

layer with the API deployment

Correct Answer: A
Section:
Explanation:
API Gateway Mocking:This feature is built for decoupling development dependencies. Here's the process:
Create resources and methods in your API Gateway.
Set the integration type to 'MOCK'.
Define Integration Responses, mapping HTTP status codes to desired mocked responses (JSON, etc.).
Deployment and Use:
Create a deployment stage for the API.
Frontend teams can call this API and get the mocked responses without a real backend.
Mocking API Gateway APIs:https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-mock-integration.html

QUESTION 39
A developer is creating an application that will give users the ability to store photos from their cellphones in the cloud. The application needs to support tens of thousands of users. The application uses an Amazon API Gateway REST API that is integrated with AWS Lambda functions to process the photos. The application stores details about the photos in Amazon DynamoDB.
Users need to create an account to access the application. In the application, users must be able to upload photos and retrieve previously uploaded photos. The photos will range in size from 300 KB to 5 MB.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Cognito user pools to manage user accounts. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. Use the Lambda function to store the photos and details in the DynamoDB table. Retrieve previously uploaded photos directly from the DynamoDB table.
B. Use Amazon Cognito user pools to manage user accounts. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
C. Create an IAM user for each user of the application during the sign-up process. Use IAM authentication to access the API Gateway API. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
D. Create a users table in DynamoDB. Use the table to manage user accounts. Create a Lambda authorizer that validates user credentials against the users table. Integrate the Lambda authorizer with API Gateway to control access to the API. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as par of the photo details in the DynamoDB table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.

Correct Answer: B
Section:
Explanation:
Amazon Cognito user pools is a service that provides a secure user directory that scales to hundreds of millions of users. The developer can use Amazon Cognito user pools to manage user accounts and create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. The developer can use the Lambda function to store the photos in Amazon S3, which is a highly scalable, durable, and secure object storage service. The developer can store the object's S3 key as part of the photo details in the DynamoDB table, which is a fast and flexible NoSQL database service. The developer can retrieve previously uploaded photos by querying DynamoDB for the S3 key and fetching the photos from S3. This solution will meet the requirements with the least operational overhead.
Reference:
[Amazon Cognito User Pools]
[Use Amazon Cognito User Pools - Amazon API Gateway]
[Amazon Simple Storage Service (S3)]
[Amazon DynamoDB]

QUESTION 40
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 partner.
Configure the Lambda function to publish messages for each partner to the partner's SNS topic.
B. Create a different Lambda function for each partner. Configure the Lambda function to notify each partner's service endpoint directly.
C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure the Lambda function to publish messages with specific attributes to the SNS topic. Subscribe each partner to the SNS topic. Apply the appropriate filter policy to the topic subscriptions.
D. Create one Amazon Simple Notification Service (Amazon SNS) topic. Subscribe all partners to the SNS topic.

Correct Answer: C
Section:
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.
Reference:
[Amazon Simple Notification Service (SNS)]
[Filtering Messages with Attributes - Amazon Simple Notification Service]

QUESTION 41
A financial company must store original customer records for 10 years for legal reasons. A complete record contains personally identifiable information (PII). According to local regulations, PII is available to only certain people in the company and must not be shared with third parties. The company needs to make the records available to third-party organizations for statistical analysis without sharing the PII.
A developer wants to store the original immutable record in Amazon S3. Depending on who accesses the S3 document, the document should be returned as is or with all the PII removed. The developer has written an AWS Lambda function to remove the PII from the document. The function is named removePii.
What should the developer do so that the company can meet the PII requirements while maintaining only one copy of the document?

A. Set up an S3 event notification that invokes the removePii function when an S3 GET request is made. Call Amazon S3 by using a GET request to access the object without PII.
B. Set up an S3 event notification that invokes the removePii function when an S3 PUT request is made. Call Amazon S3 by using a PUT request to access the object without PII.
C. Create an S3 Object Lambda access point from the S3 console. Select the removePii function. Use S3 Access Points to access the object without PII.
D. Create an S3 access point from the S3 console. Use the access point name to call the GetObjectLegalHold S3 API function. Pass in the removePii function name to access the object without PII.

Correct Answer: C
Section:
Explanation:
S3 Object Lambda allows you to add your own code to process data retrieved from S3 before returning it to an application. You can use an AWS Lambda function to modify the data, such as removing PII, redacting confidential information, or resizing images. You can create an S3 Object Lambda access point and associate it with your Lambda function. Then, you can use the access point to request objects from S3 and get the modified data back. This way, you can maintain only one copy of the original document in S3 and apply different transformations depending on who accesses it.
Reference: Using AWS Lambda with Amazon S3

QUESTION 42
A developer is deploying an AWS Lambda function The developer wants the ability to return to older versions of the function quickly and seamlessly.
How can the developer achieve this goal with the LEAST operational overhead?

A. Use AWS OpsWorks to perform blue/green deployments.
B. Use a function alias with different versions.
C. Maintain deployment packages for older versions in Amazon S3.
D. Use AWS CodePipeline for deployments and rollbacks.

Correct Answer: B

Section:
Explanation:
A function alias is a pointer to a specific Lambda function version. You can use aliases to create different environments for your function, such as development, testing, and production. You can also use aliases to perform blue/green deployments by shifting traffic between two versions of your function gradually. This way, you can easily roll back to a previous version if something goes wrong, without having to redeploy your code or change your configuration. Reference: AWS Lambda function aliases

QUESTION 43
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.

Correct Answer: B
Section:
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

QUESTION 44
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

Correct Answer: B
Section:
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]

QUESTION 45
A company is building a serverless application on AWS. The application uses an AWS Lambda function to process customer orders 24 hours a day, 7 days a week. The Lambda function calls an external vendor's HTTP API to process payments.
During load tests, a developer discovers that the external vendor payment processing API occasionally times out and returns errors. The company expects that some payment processing API calls will return errors.
The company wants the support team to receive notifications in near real time only when the payment processing external API error rate exceed 5% of the total number of transactions in an hour. Developers need to use an existing Amazon Simple Notification Service (Amazon SNS) topic that is configured to notify the support team.
Which solution will meet these requirements?

A. Write the results of payment processing API calls to Amazon CloudWatch. Use Amazon CloudWatch Logs Insights to query the CloudWatch logs. Schedule the Lambda function to check the CloudWatch logs and notify the existing SNS topic.
B. Publish custom metrics to CloudWatch that record the failures of the external payment processing API calls. Configure a CloudWatch alarm to notify the existing SNS topic when error rate exceeds the specified rate.
C. Publish the results of the external payment processing API calls to a new Amazon SNS topic. Subscribe the support team members to the new SNS topic.
D. Write the results of the external payment processing API calls to Amazon S3. Schedule an Amazon Athena query to run at regular intervals. Configure Athena to send notifications to the existing SNS topic when the error rate exceeds the specified rate.

Correct Answer: B
Section:
Explanation:
Amazon CloudWatch is a service that monitors AWS resources and applications. The developer can publish custom metrics to CloudWatch that record the failures of the external payment processing API calls. The developer can configure a CloudWatch alarm to notify the existing SNS topic when the error rate exceeds 5% of the total number of transactions in an hour. This solution will meet the requirements in a near real-time and scalable way.
Reference:
[What Is Amazon CloudWatch? - Amazon CloudWatch]
[Publishing Custom Metrics - Amazon CloudWatch]
[Creating Amazon CloudWatch Alarms - Amazon CloudWatch]

QUESTION 46
A company is offering APIs as a service over the internet to provide unauthenticated read access to statistical information that is updated daily. The company uses Amazon API Gateway and AWS Lambda to develop the APIs. The service has become popular, and the company wants to enhance the responsiveness of the APIs.
Which action can help the company achieve this goal?

A. Enable API caching in API Gateway.
B. Configure API Gateway to use an interface VPC endpoint.
C. Enable cross-origin resource sharing (CORS) for the APIs.
D. Configure usage plans and API keys in API Gateway.

Correct Answer: A
Section:
Explanation:
Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. The developer can enable API caching in API Gateway to cache responses from the backend integration point for a specified time-to-live (TTL) period. This can improve the responsiveness of the APIs by reducing the number of calls made to the backend service.
Reference:
[What Is Amazon API Gateway? - Amazon API Gateway]
[Enable API Caching to Enhance Responsiveness - Amazon API Gateway]

QUESTION 47
A developer wants to store information about movies. Each movie has a title, release year, and genre. The movie information also can include additional properties about the cast and production crew. This additional information is inconsistent across movies. For example, one movie might have an assistant director, and another movie might have an animal trainer.
The developer needs to implement a solution to support the following use cases:
For a given title and release year, get all details about the movie that has that title and release year.
For a given title, get all details about all movies that have that title.
For a given genre, get all details about all movies in that genre.
Which data store configuration will meet these requirements?

A. Create an Amazon DynamoDB table. Configure the table with a primary key that consists of the title as the partition key and the release year as the sort key. Create a global secondary index that uses the genre as the partition key and the title as the sort key.

B. Create an Amazon DynamoDB table. Configure the table with a primary key that consists of the genre as the partition key and the release year as the sort key. Create a global secondary index that uses the title as the partition key.
C. On an Amazon RDS DB instance, create a table that contains columns for title, release year, and genre. Configure the title as the primary key.
D. On an Amazon RDS DB instance, create a table where the primary key is the title and all other data is encoded into JSON format as one additional column.

Correct Answer: A
Section:
Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance with seamless scalability. The developer can create a DynamoDB table and configure the table with a primary key that consists of the title as the partition key and the release year as the sort key. This will enable querying for a given title and release year efficiently. The developer can also create a global secondary index that uses the genre as the partition key and the title as the sort key.
This will enable querying for a given genre efficiently. The developer can store additional properties about the cast and production crew as attributes in the DynamoDB table. These attributes can have different data types and structures, and they do not need to be consistent across items.
Reference:
[Amazon DynamoDB]
[Working with Queries - Amazon DynamoDB]
[Working with Global Secondary Indexes - Amazon DynamoDB]

QUESTION 48
A developer maintains an Amazon API Gateway REST API. Customers use the API through a frontend UI and Amazon Cognito authentication.
The developer has a new version of the API that contains new endpoints and backward-incompatible interface changes. The developer needs to provide beta access to other developers on the team without affecting customers.
Which solution will meet these requirements with the LEAST operational overhead?

A. Define a development stage on the API Gateway API. Instruct the other developers to point the endpoints to the development stage.
B. Define a new API Gateway API that points to the new API application code. Instruct the other developers to point the endpoints to the new API.
C. Implement a query parameter in the API application code that determines which code version to call.
D. Specify new API Gateway endpoints for the API endpoints that the developer wants to add.

Correct Answer: A
Section:
Explanation:
Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. The developer can define a development stage on the API Gateway API and instruct the other developers to point the endpoints to the development stage. This way, the developer can provide beta access to the new version of the API without affecting customers who use the production stage. This solution will meet the requirements with the least operational overhead.
Reference:
[What Is Amazon API Gateway? - Amazon API Gateway]
[Set up a Stage in API Gateway - Amazon API Gateway]

QUESTION 49
A developer is creating an application that will store personal health information (PHI). The PHI needs to be encrypted at all times. An encrypted Amazon RDS for MySQL DB instance is storing the dat a. The developer wants to increase the performance of the application by caching frequently accessed data while adding the ability to sort or rank the cached datasets.
Which solution will meet these requirements?

A. Create an Amazon ElastiCache for Redis instance. Enable encryption of data in transit and at rest. Store frequently accessed data in the cache.
B. Create an Amazon ElastiCache for Memcached instance. Enable encryption of data in transit and at rest. Store frequently accessed data in the cache.
C. Create an Amazon RDS for MySQL read replica. Connect to the read replica by using SSL. Configure the read replica to store frequently accessed data.
D. Create an Amazon DynamoDB table and a DynamoDB Accelerator (DAX) cluster for the table. Store frequently accessed data in the DynamoDB table.

Correct Answer: A
Section:
Explanation:
Amazon ElastiCache is a service that offers fully managed in-memory data stores that are compatible with Redis or Memcached. The developer can create an ElastiCache for Redis instance and enable encryption of data in transit and at rest. This will ensure that the PHI is encrypted at all times. The developer can store frequently accessed data in the cache and use Redis features such as sorting and ranking to enhance the performance of the application.
Reference:
[What Is Amazon ElastiCache? - Amazon ElastiCache]
[Encryption in Transit - Amazon ElastiCache for Redis]
[Encryption at Rest - Amazon ElastiCache for Redis]

QUESTION 50
A company has a multi-node Windows legacy application that runs on premises. The application uses a network shared folder as a centralized configuration repository to store configuration files in .xml format. The company is migrating the application to Amazon EC2 instances. As part of the migration to AWS, a developer must identify a solution that provides high availability for the repository.
Which solution will meet this requirement MOST cost-effectively?

A. Mount an Amazon Elastic Block Store (Amazon EBS) volume onto one of the EC2 instances. Deploy a file system on the EBS volume. Use the host operating system to share a folder. Update the application code to read and write configuration files from the shared folder.
B. Deploy a micro EC2 instance with an instance store volume. Use the host operating system to share a folder. Update the application code to read and write configuration files from the shared folder.
C. Create an Amazon S3 bucket to host the repository. Migrate the existing .xml files to the S3 bucket. Update the application code to use the AWS SDK to read and write configuration files from Amazon S3.
D. Create an Amazon S3 bucket to host the repository. Migrate the existing .xml files to the S3 bucket. Mount the S3 bucket to the EC2 instances as a local volume. Update the application code to read and write configuration files from the disk.

Correct Answer: C
Section:
Explanation:
Amazon S3 is a service that provides highly scalable, durable, and secure object storage. The developer can create an S3 bucket to host the repository and migrate the existing .xml files to the S3 bucket. The developer can update the application code to use the AWS SDK to read and write configuration files from S3. This solution will meet the requirement of high availability for the repository in a cost-effective way.
Reference:
[Amazon Simple Storage Service (S3)] [Using AWS SDKs with Amazon S3]

QUESTION 51
A company wants to deploy and maintain static websites on AWS. Each website's source code is hosted in one of several version control systems, including AWS CodeCommit, Bitbucket, and GitHub.
The company wants to implement phased releases by using development, staging, user acceptance testing, and production environments in the AWS Cloud. Deployments to each environment must be started by code merges on the relevant Git branch. The company wants to use HTTPS for all data exchange. The company needs a solution that does not require servers to run continuously.
Which solution will meet these requirements with the LEAST operational overhead?

A. Host each website by using AWS Amplify with a serverless backend. Conned the repository branches that correspond to each of the desired environments. Start deployments by merging code changes to a desired branch.
B. Host each website in AWS Elastic Beanstalk with multiple environments. Use the EB CLI to link each repository branch. Integrate AWS CodePipeline to automate deployments from version control code merges.
C. Host each website in different Amazon S3 buckets for each environment. Configure AWS CodePipeline to pull source code from version control. Add an AWS CodeBuild stage to copy source code to Amazon S3.
D. Host each website on its own Amazon EC2 instance. Write a custom deployment script to bundle each website's static assets. Copy the assets to Amazon EC2. Set up a workflow to run the script when code is merged.

Correct Answer: A
Section:
Explanation:
AWS Amplify is a set of tools and services that enables developers to build and deploy full-stack web and mobile applications that are powered by AWS. AWS Amplify supports hosting static websites on Amazon S3 and Amazon CloudFront, with HTTPS enabled by default. AWS Amplify also integrates with various version control systems, such as AWS CodeCommit, Bitbucket, and GitHub, and allows developers to connect different branches to different environments. AWS Amplify automatically builds and deploys the website whenever code changes are merged to a connected branch, enabling phased releases with minimal operational overhead. Reference: AWS Amplify Console

QUESTION 52
A company is migrating an on-premises database to Amazon RDS for MySQL. The company has readheavy workloads. The company wants to refactor the code to achieve optimum read performance for queries.
Which solution will meet this requirement with LEAST current and future effort?

A. Use a multi-AZ Amazon RDS deployment. Increase the number of connections that the code makes to the database or increase the connection pool size if a connection pool is in use.
B. Use a multi-AZ Amazon RDS deployment. Modify the code so that queries access the secondary RDS instance.
C. Deploy Amazon RDS with one or more read replicas. Modify the application code so that queries use the URL for the read replicas.
D. Use open source replication software to create a copy of the MySQL database on an Amazon EC2 instance. Modify the application code so that queries use the IP address of the EC2 instance.

Correct Answer: C
Section:
Explanation:
Amazon RDS for MySQL supports read replicas, which are copies of the primary database instance that can handle read-only queries. Read replicas can improve the read performance of the database by offloading the read workload from the primary instance and distributing it across multiple replicas. To use read replicas, the application code needs to be modified to direct read queries to the URL of the read replicas, while write queries still go to the URL of the primary instance. This solution requires less current and future effort than using a multi-AZ deployment, which does not provide read scaling benefits, or using open source replication software, which requires additional configuration and maintenance. Reference: Working with read replicas

QUESTION 53
An application that is hosted on an Amazon EC2 instance needs access to files that are stored in an Amazon S3 bucket. The application lists the objects that are stored in the S3 bucket and displays a table to the user. During testing, a developer discovers that the application does not show any objects in the list.
What is the MOST secure way to resolve this issue?

A. Update the IAM instance profile that is attached to the EC2 instance to include the S3:* permission for the S3 bucket.
B. Update the IAM instance profile that is attached to the EC2 instance to include the S3:ListBucket permission for the S3 bucket.
C. Update the developer's user permissions to include the S3:ListBucket permission for the S3 bucket.
D. Update the S3 bucket policy by including the S3:ListBucket permission and by setting the Principal element to specify the account number of the EC2 instance.

Correct Answer: B
Section:
Explanation:
IAM instance profiles are containers for IAM roles that can be associated with EC2 instances. An IAM role is a set of permissions that grant access to AWS resources. An IAM role can be used to allow an EC2 instance to access an S3 bucket by including the appropriate permissions in the role's policy. The S3:ListBucket permission allows listing the objects in an S3 bucket. By updating the IAM instance profile with this permission, the application on the EC2 instance can retrieve the objects from the S3 bucket and display them to the user. Reference: Using an IAM role to grant permissions to applications running on Amazon EC2 instances

QUESTION 54
A company is planning to securely manage one-time fixed license keys in AWS. The company's development team needs to access the license keys in automaton scripts that run in Amazon EC2 instances and in AWS CloudFormation stacks.
Which solution will meet these requirements MOST cost-effectively?

A. Amazon S3 with encrypted files prefixed with "config"
B. AWS Secrets Manager secrets with a tag that is named SecretString
C. AWS Systems Manager Parameter Store SecureString parameters
D. CloudFormation NoEcho parameters

Correct Answer: C
Section:
Explanation:
AWS Systems Manager Parameter Store is a service that provides secure, hierarchical storage for configuration data and secrets. Parameter Store supports SecureString parameters, which are encrypted using AWS Key

Management Service (AWS KMS) keys. SecureString parameters can be used to store license keys in AWS and retrieve them securely from automation scripts that run in EC2 instances or CloudFormation stacks. Parameter Store is a cost-effective solution because it does not charge for storing parameters or API calls. Reference: Working with Systems Manager parameters

QUESTION 55
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 function. Configure VPC1 access for the function. Attach a security group named SG1 to both the Lambda function and the database. Configure the security group inbound and outbound rules to allow TCP traffic on Port 3306.
B. 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.
C. Create the Lambda function. Configure VPC1 access for the function. Assign a security group named SG1 to the Lambda function. Assign a second security group named SG2 to the database. Add an inbound rule to SG1 to allow TCP traffic from Port 3306.
D. 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.

Correct Answer: A
Section:
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]

QUESTION 56
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

Correct Answer: B, D
Section:
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.
Reference:
[What Is Amazon API Gateway? - Amazon API Gateway]
[Amazon API Gateway Metrics and Dimensions - Amazon CloudWatch]
[Troubleshooting API Errors - Amazon API Gateway]

QUESTION 57
A developer is writing a serverless application that requires an AWS Lambda function to be invoked every 10 minutes.
What is an automated and serverless way to invoke the function?

A. Deploy an Amazon EC2 instance based on Linux, and edit its /etc/confab file by adding a command to periodically invoke the lambda function
B. Configure an environment variable named PERIOD for the Lambda function. Set the value to 600.
C. Create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function.
D. Create an Amazon Simple Notification Service (Amazon SNS) topic that has a subscription to the Lambda function with a 600-second timer.

Correct Answer: C
Section:
Explanation:
The solution that will meet the requirements is to create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function. This way, the developer can use an automated and serverless way to invoke the function every 10 minutes. The developer can also use a cron expression or a rate expression to specify the schedule for the rule. The other options either involve using an Amazon EC2 instance, which is not serverless, or using environment variables or query parameters, which do not trigger the function.
Reference: Schedule AWS Lambda functions using EventBridge

QUESTION 58
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

Correct Answer: A
Section:
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)]

QUESTION 59
An application that is deployed to Amazon EC2 is using Amazon DynamoDB. The app cation calls the DynamoDB REST API Periodically the application receives a ProvisionedThroughputExceededException error when the application writes to a DynamoDB table.
Which solutions will mitigate this error MOST cost-effectively^ (Select TWO)

A. Modify the application code to perform exponential back off when the error is received.
B. Modify the application to use the AWS SDKs for DynamoDB.
C. Increase the read and write throughput of the DynamoDB table.
D. Create a DynamoDB Accelerator (DAX) cluster for the DynamoDB table.
E. Create a second DynamoDB table Distribute the reads and writes between the two tables.

Correct Answer: A, B
Section:
Explanation:
These solutions will mitigate the error most cost-effectively because they do not require increasing the provisioned throughput of the DynamoDB table or creating additional resources. Exponential backoff is a retry strategy that increases the waiting time between retries to reduce the number of requests sent to DynamoDB. The AWS SDKs for DynamoDB implement exponential backoff by default and also provide other features such as automatic pagination and encryption. Increasing the read and write throughput of the DynamoDB table, creating a DynamoDB Accelerator (DAX) cluster, or creating a second DynamoDB table will incur additional costs and complexity.
Reference: [Error Retries and Exponential Backoff in AWS], [Using the AWS SDKs with DynamoDB]

QUESTION 60
When a developer tries to run an AWS Code Build project, it raises an error because the length of all environment variables exceeds the limit for the combined maximum of characters.
What is the recommended solution?

A. Add the export LC-_ALL" on _ US, tuft" command to the pre _ build section to ensure POSIX Localization.
B. Use Amazon Cognate to store key-value pairs for large numbers of environment variables
C. Update the settings for the build project to use an Amazon S3 bucket for large numbers of environment variables
D. Use AWS Systems Manager Parameter Store to store large numbers ot environment variables

Correct Answer: D
Section:
Explanation:
This solution allows the developer to overcome the limit for the combined maximum of characters for environment variables in AWS CodeBuild. AWS Systems Manager Parameter Store provides secure, hierarchical storage for configuration data management and secrets management. The developer can store large numbers of environment variables as parameters in Parameter Store and reference them in the buildspec file using parameter references. Adding export LC_ALL="en_US.utf8" command to the pre_build section will not affect the environment variables limit. Using Amazon Cognito or an Amazon S3 bucket to store key-value pairs for environment variables will require additional configuration and integration.
Reference: [Build Specification Reference for AWS CodeBuild], [What Is AWS Systems Manager Parameter Store?]

QUESTION 61
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 policy. Store the access key ID and secret access key on the EC2 instances
C. Modify the application to use the S3 GeneratePresignedUrl API call
D. Modify the application to use the S3 GetObject API call and to return the object handle to the user
E. Modify the application to delegate requests to the S3 bucket.

Correct Answer: A, C
Section:
Explanation:
IAM Roles for EC2 (A):The most secure way to provide AWS permissions from EC2.
Create a role with a policy allowings3:GetObjecton the specific bucket.
Attach the role to an instance profile and associate that profile with your instances.
Pre-signed URLs (C):Temporary, authenticated URLs for specific S3 actions.
Modify the app to use the AWS SDK to callGeneratePresignedUrl.
Embed these URLs when a user is properly logged in, allowing download access.
IAM Roles for EC2:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html
Generating Presigned URLs:https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.htm

QUESTION 62
An AWS Lambda function requires read access to an Amazon S3 bucket and requires read/write access to an Amazon DynamoDB table The correct 1AM policy already exists
What is the MOST secure way to grant the Lambda function access to the S3 bucket and the DynamoDB table?

A. Attach the existing 1AM policy to the Lambda function.
B. Create an 1AM role for the Lambda function Attach the existing 1AM policy to the role Attach the role to the Lambda function
C. Create an 1AM user with programmatic access Attach the existing 1AM policy to the user. Add the user access key ID and secret access key as environment variables in the Lambda function.
D. Add the AWS account root user access key ID and secret access key as encrypted environment variables in the Lambda function

Correct Answer: B
Section:
Explanation:
Principle of Least Privilege:Granting specific permissions through an IAM role is more secure than directly attaching policies to a function or using root user credentials.
IAM Roles for Lambda:Designed to provide temporary credentials to Lambda functions, enhancing security.
Reusability:The existing IAM policy ensures the correct S3 and DynamoDB access is granted.
IAM Roles for Lambda Documentation:https://docs.aws.amazon.com/lambda/latest/dg/lambda-intro-execution-role.html
IAM Best Practices:https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html

QUESTION 63
A developer is designing a serverless application for a game in which users register and log in through a web browser The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API
The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.
Which solution will meet these requirements'?

A. Create Amazon Cognito user pools for external social identity providers Configure 1AM roles for the identity pools.
B. Program the sign-in page to create users' 1AM groups with the 1AM roles attached to the groups
C. Create an Amazon RDS for SQL Server DB instance to store the users and manage the permissions to the backend resources in AWS
D. Configure the sign-in page to register and store the users and their passwords in an Amazon DynamoDB table with an attached IAM policy.

Correct Answer: A
Section:
Explanation:
Amazon Cognito User Pools:A managed user directory service, simplifying user registration and login.
Social Identity Providers:Cognito supports integration with external providers (e.g., Google, Facebook), reducing development effort.
IAM Roles for Authorization:Cognito-managed IAM roles grant fine-grained access to AWS resources (like Lambda functions).
Operational Overhead:Cognito minimizes the need to manage user identities and credentials independently.
Amazon Cognito Documentationhttps://docs.aws.amazon.com/cognito/
Cognito User Pools for Web Applications:https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-user-pools-app-integration.html

QUESTION 64
A developer supports an application that accesses data in an Amazon DynamoDB table. One of the item attributes is expirationDate in the timestamp format. The application uses this attribute to find items, archive them, and remove them from the table based on the timestamp value
The application will be decommissioned soon, and the developer must find another way to implement this functionality. The developer needs a solution that will require the least amount of code to write.
Which solution will meet these requirements?

A. Enable TTL on the expirationDate attribute in the table. Create a DynamoDB stream. Create an AWS Lambda function to process the deleted items. Create a DynamoDB trigger for the Lambda function.
B. Create two AWS Lambda functions one to delete the items and one to process the items Create a DynamoDB stream Use the DeleteItem API operation to delete the items based on the expirationDate attribute Use the GetRecords API operation to get the items from the DynamoDB stream and process them
C. Create two AWS Lambda functions, one to delete the items and one to process the items. Create an Amazon EventBndge scheduled rule to invoke the Lambda Functions Use the DeleteItem API operation to delete the items based on the expirationDate attribute. Use the GetRecords API operation to get the items from the DynamoDB table and process them.
D. Enable TTL on the expirationDate attribute in the table Specify an Amazon Simple Queue Service (Amazon SQS> dead-letter queue as the target to delete the items Create an AWS Lambda function to process the items

Correct Answer: A
Section:
Explanation:
TTL for Automatic Deletion:DynamoDB's Time-to-Live effortlessly deletes expired items without manual intervention.
DynamoDB Stream:Captures changes to the table, including deletions of expired items, triggering downstream actions.
Lambda for Processing:A Lambda function connected to the stream provides custom logic for handling the deleted items.
Code Efficiency:This solution leverages native DynamoDB features and stream-based processing, minimizing the need for custom code.
DynamoDB TTL Documentation:https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html
DynamoDB Streams Documentation:https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html

QUESTION 65
A developer uses AWS CloudFormation to deploy an Amazon API Gateway API and an AWS Step Functions state machine The state machine must reference the API Gateway API after the CloudFormation template is deployed
The developer needs a solution that uses the state machine to reference the API Gateway endpoint.
Which solution will meet these requirements MOST cost-effectively?

A. Configure the CloudFormation template to reference the API endpoint in the DefinitionSubstitutions property for the AWS StepFunctions StateMachme resource.
B. Configure the CloudFormation template to store the API endpoint in an environment variable for the AWS::StepFunctions::StateMachine resourc Configure the state machine to reference the environment variable

C. Configure the CloudFormation template to store the API endpoint in a standard AWS: SecretsManager Secret resource Configure the state machine to reference the resource
D. Configure the CloudFormation template to store the API endpoint in a standard AWS::AppConfig;:ConfigurationProfile resource Configure the state machine to reference the resource.

Correct Answer: A
Section:
Explanation:
CloudFormation and Dynamic
Reference:TheDefinitionSubstitutionsproperty in CloudFormation allows you to pass values into Step Functions state machines at runtime.
Cost-Effectiveness:This solution is cost-effective as it leverages CloudFormation's built-in capabilities, avoiding the need for additional services like Secrets Manager or AppConfig.
AWS Step Functions State Machine:https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-stepfunctions-statemachine.html
CloudFormation DefinitionSubstitutions:https://github.com/aws-cloudformation/aws-cloudformation-resource-providers-stepfunctions/issues/14

QUESTION 66
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. Enable AWS X-Ray active tracing in the Lambda function Review the logs in X-Ray
B. Configure AWS CloudTrail View the trail logs that are associated with the Lambda function.
C. Review the AWS Config logs in Amazon Cloud Watch.
D. Review the Amazon CloudWatch logs that are associated with the Lambda function.

Correct Answer: A
Section:
Explanation:
Tracing Distributed Systems:AWS X-Ray is designed to trace requests across services, helping identify bottlenecks in distributed applications like this one.
No Code Changes:Enabling X-Ray tracing often requires minimal code changes, meeting the requirement.
Identifying Bottlenecks: Analyzing X-Ray traces and logs will reveal latency in communications between different AWS services, leading to the high duration time.
AWS X-Ray:https://aws.amazon.com/xray/
X-Ray and Lambda:https://docs.aws.amazon.com/xray/latest/devguide/xray-services-lambda.html

QUESTION 67
A developer designed an application on an Amazon EC2 instance The application makes API requests to objects in an Amazon S3 bucket
Which combination of steps will ensure that the application makes the API requests in the MOST secure manner? (Select TWO.)

A. Create an IAM user that has permissions to the S3 bucket. Add the user to an 1AM group
B. Create an IAM role that has permissions to the S3 bucket
C. Add the IAM role to an instance profile. Attach the instance profile to the EC2 instance.
D. Create an 1AM role that has permissions to the S3 bucket Assign the role to an 1AM group
E. Store the credentials of the IAM user in the environment variables on the EC2 instance

Correct Answer: B, C
Section:
Explanation:
IAM Roles for EC2: IAM roles are the recommended way to provide AWS credentials to applications running on EC2 instances. Here's how this works:
You create an IAM role with the necessary permissions to access the target S3 bucket.
You create an instance profile and associate the IAM role with this profile.
When launching the EC2 instance, you attach this instance profile.

Temporary Security Credentials: When the application on the EC2 instance needs to access S3, it doesn't directly use access keys. Instead, the AWS SDK running on the instance retrieves temporary security credentials associated with the role. These are rotated automatically by AWS.
IAM Roles for Amazon EC2:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html
Temporary Security Credentials:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html

QUESTION 68
A developer is working on an ecommerce website The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available
How can the developer update the application to meet these requirements with MINIMUM changes?

A. Rewrite the application to be cloud native and to run on AWS Lambda, where the logs can be reviewed in Amazon CloudWatch
B. Set up centralized logging by using Amazon OpenSearch Service, Logstash, and OpenSearch Dashboards
C. Scale down the application to one larger EC2 instance where only one instance is recording logs
D. Install the unified Amazon CloudWatch agent on the EC2 instances Configure the agent to push the application logs to CloudWatch

Correct Answer: D
Section:
Explanation:
Centralized Logging Benefits: Centralized logging is essential for operational visibility in scalable systems, especially those using multiple EC2 instances like our e-commerce website. CloudWatch provides this capability, along with other monitoring features.
CloudWatch Agent: This is the best way to send custom application logs from EC2 instances to CloudWatch. Here's the process:
Install the CloudWatch agent on each EC2 instance.
Configure the agent with a configuration file, specifying:
Which log files to collect.
The format in which to send logs to CloudWatch (e.g., JSON).
The specific CloudWatch Logs log group and log stream for these logs.
Viewing and Analyzing Logs: Once the agent is pushing logs, use the CloudWatch Logs console or API:
View and search the logs across all instances.
Set up alarms based on log events.
Use CloudWatch Logs Insights for sophisticated queries and analysis.
Amazon CloudWatch Logs:https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html
Unified CloudWatch Agent:https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AgentReference.html
CloudWatch Logs Insights:https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.html

QUESTION 69
A company has an existing application that has hardcoded database credentials A developer needs to modify the existing application The application is deployed in two AWS Regions with an active-passive failover configuration to meet company's disaster recovery strategy
The developer needs a solution to store the credentials outside the code. The solution must comply With the company's disaster recovery strategy
Which solution Will meet these requirements in the MOST secure way?

A. Store the credentials in AWS Secrets Manager in the primary Region. Enable secret replication to the secondary Region Update the application to use the Amazon Resource Name (ARN) based on the Region.
B. Store credentials in AWS Systems Manager Parameter Store in the primary Region. Enable parameter replication to the secondary Region. Update the application to use the Amazon Resource Name (ARN) based on the Region.
C. Store credentials in a config file. Upload the config file to an S3 bucket in me primary Region. Enable Cross-Region Replication (CRR) to an S3 bucket in the secondary region. Update the application to access the config file from the S3 bucket based on the Region.
D. Store credentials in a config file. Upload the config file to an Amazon Elastic File System (Amazon EFS) file system. Update the application to use the Amazon EFS file system Regional endpoints to access the config file in the primary and secondary Regions.

Correct Answer: A

Section:
Explanation:
AWS Secrets Manager is a service that allows you to store and manage secrets, such as database credentials, API keys, and passwords, in a secure and centralized way.It also provides features such as automatic secret rotation, auditing, and monitoring1. By using AWS Secrets Manager, you can avoid hardcoding credentials in your code, which is a bad security practice and makes it difficult to update them.You can also replicate your secrets to another Region, which is useful for disaster recovery purposes2. To access your secrets from your application, you can use the ARN of the secret, which is a unique identifier that includes the Region name.This way, your application can use the appropriate secret based on the Region where it is deployed3.
AWS Secrets Manager
Replicating and sharing secrets
Using your own encryption keys

QUESTION 70
A developer is creating an AWS Lambda function that searches for items from an Amazon DynamoDB table that contains customer contact information- The DynamoDB table items have the customer's email_address as the partition key and additional properties such as customer_type, name, and job_tltle.
The Lambda function runs whenever a user types a new character into the customer_type text input The developer wants the search to return partial matches of all the email_address property of a particular customer_type The developer does not want to recreate the DynamoDB table.
What should the developer do to meet these requirements?

A. Add a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key Perform a query operation on the GSI by using the begvns_wth key condition expression With the emad_address property
B. Add a global secondary index (GSI) to the DynamoDB table With ernail_address as the partition key and customer_type as the sort key Perform a query operation on the GSI by using the begins_wtth key condition expression With the emal_address property.
C. Add a local secondary index (LSI) to the DynamoDB table With customer_type as the partition key and email_address as the sort key Perform a query operation on the LSI by using the begins_wlth key condition expression With the email_address property
D. Add a local secondary Index (LSI) to the DynamoDB table With job_tltle as the partition key and emad_address as the sort key Perform a query operation on the LSI by using the begins_wrth key condition expression With the email_address property

Correct Answer: A
Section:
Explanation:
Understand the Problem: The existing DynamoDB table has email_address as the partition key. Searching by customer_type requires a different data access pattern. We need an efficient way to query for partial matches on email_address based on customer_type.
Why Global Secondary Index (GSI):
GSIs allow you to define a different partition key and sort key from the main table, enabling new query patterns.
In this case, havingcustomer_typeas the GSI's partition key lets you group all emails with the same customer type together.
Usingemail_addressas the sort key allows ordering within each customer type, facilitating the partial matching.
Querying the GSI:
You'll perform a query operation on the GSI, not the original table.
Use thebegins_withkey condition expression on the GSI's sort key (email_address) to find partial matches as the user types in thecustomer_typefield.
DynamoDB Global Secondary Indexes:https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html
DynamoDB Query Operation:[invalid URL removed]

QUESTION 71
A developer is deploying a company's application to Amazon EC2 instances The application generates gigabytes of data files each day The files are rarely accessed but the files must be available to the application's users within minutes of a request during the first year of storage The company must retain the files for 7 years.
How can the developer implement the application to meet these requirements MOST cost-effectively?

A. Store the files in an Amazon S3 bucket Use the S3 Glacier Instant Retrieval storage class Create an S3 Lifecycle policy to transition the files to the S3 Glacier Deep Archive storage class after 1 year
B. Store the files in an Amazon S3 bucket. Use the S3 Standard storage class. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Flexible Retrieval storage class after 1 year.
C. Store the files on an Amazon Elastic Block Store (Amazon EBS) volume Use Amazon Data Lifecycle Manager (Amazon DLM) to create snapshots of the EBS volumes and to store those snapshots in Amazon S3

D. Store the files on an Amazon Elastic File System (Amazon EFS) mount. Configure EFS lifecycle management to transition the files to the EFS Standard-Infrequent Access (Standard-IA) storage class after 1 year.

Correct Answer: A
Section:
Explanation:
Amazon S3 Glacier Instant Retrieval is an archive storage class that delivers the lowest-cost storage for long-lived data that is rarely accessed and requires retrieval in milliseconds. With S3 Glacier Instant Retrieval, you can save up to 68% on storage costs compared to using the S3 Standard-Infrequent Access (S3 Standard-IA) storage class, when your data is accessed once per quarter. https://aws.amazon.com/s3/storage-classes/glacier/instant-retrieval/
Understanding Storage Requirements:
Files are large and infrequently accessed, but need to be available within minutes when requested in the first year.
Long-term (7-year) retention is required.
Cost-effectiveness is a top priority.
Why S3 Glacier Instant Retrieval:
Matches the retrieval requirements (access within minutes).
More cost-effective than S3 Standard for infrequently accessed data.
Simpler to use than traditional Glacier where retrievals take hours.
Why S3 Glacier Deep Archive:
Most cost-effective S3 storage class for long term archival.
Meets the 7-year retention requirement.
S3 Lifecycle Policy:
Automate the transition from Glacier Instant Retrieval to Glacier Deep Archive after one year.
Optimize costs by matching storage classes to access patterns.
Amazon S3 Storage Classes:https://aws.amazon.com/s3/storage-classes/
S3 Glacier Instant Retrieval:[invalid URL removed]
S3 Glacier Deep Archive:[invalid URL removed]
S3 Lifecycle Policies:https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html

QUESTION 72
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

Correct Answer: D
Section:
Explanation:
CloudFormation and Lambda Environment Variables:
CloudFormation is an excellent tool to manage infrastructure as code, including the log group resource.
Lambda functions can access environment variables at runtime, making them a suitable way to pass configuration information like the log group ARN.
CloudFormation Template Modification:
In your CloudFormation template, define the log group resource.
In the Lambda function resource, add anEnvironmentsection:
YAML
Environment:

Variables:
LOG_GROUP_ARN: !Ref LogGroupResourceName
Use codewith caution.
content_copy
The!Refintrinsic function retrieves the log group's ARN, which CloudFormation generates during stack creation.
Using the ARN in Your Lambda Function:
Within your Lambda code, access theLOG_GROUP_ARNenvironment variable.
Configure your logging library (e.g., Python'sloggingmodule) to send logs to the specified log group.
AWS Lambda Environment Variables:https://docs.aws.amazon.com/lambda/latest/dg/configuration-envvars.html
CloudFormation !Ref Intrinsic Function:https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-ref.html

QUESTION 73
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 AMI. Relaunch the stack for both Regions.
B. 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
C. 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
D. Manually deploy the application outside AWS CloudFormation in us-west-1.

Correct Answer: B
Section:
Explanation:
Problem: CloudFormation can't find the custom AMI in the target region (us-west-1) because AMIs are region-specific.
Copying AMIs:
AMIs can be copied across regions, maintaining their configuration.
This approach minimizes operational overhead as the existing CloudFormation template can be reused with a minor update.
Updating the Template:
Modify the CloudFormation template in us-west-1 to reference the newly copied AMI's ID in that region.
Copying AMIs:https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/CopyingAMIs.html
CloudFormation Templates and AMIs:[invalid URL removed]

QUESTION 74
A company is creating an application that processes csv files from Amazon S3 A developer has created an S3 bucket The developer has also created an AWS Lambda function to process the csv files from the S3 bucket
Which combination of steps will invoke the Lambda function when a csv file is uploaded to Amazon S3? (Select TWO.)

A. Create an Amazon EventBridge rule Configure the rule with a pattern to match the S3 object created event
B. Schedule an Amazon EventBridge rule to run a new Lambda function to scan the S3 bucket.
C. Add a trigger to the existing Lambda function. Set the trigger type to EventBridge Select the Amazon EventBridge rule.
D. Create a new Lambda function to scan the S3 bucket for recently added S3 objects
E. Add S3 Lifecycle rules to invoke the existing Lambda function

Correct Answer: A, E
Section:
Explanation:

Amazon EventBridge:A service that reacts to events from various AWS sources, including S3. Rules define which events trigger actions (like invoking Lambda functions).
S3 Object Created Events:EventBridge can detect these, providing seamless integration for automated CSV processing.
S3 Lifecycle Rules:Allow for actions based on object age or prefixes. These can directly trigger Lambda functions for file processing.
Amazon EventBridge Documentation:https://docs.aws.amazon.com/eventbridge/
Working with S3 Event Notifications:https://docs.aws.amazon.com/AmazonS3/latest/userguide/EventNotifications.html
S3 Lifecycle Configuration:https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html

QUESTION 75
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 system. Mount the EFS file system in Lambda. Store the result files and log file in the mount point. Append the log entries to the log file.
B. Create an Amazon Elastic Block Store (Amazon EBS) Multi-Attach enabled volume Attach the EBS volume to all Lambda functions. Update the Lambda function code to download the log file, append the log entries, and upload the modified log file to Amazon EBS
C. Create a reference to the /tmp local directory. Store the result files and log file by using the directory reference. Append the log entry to the log file.
D. 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

Correct Answer: A
Section:
Explanation:
Amazon EFS:A network file system (NFS) providing shared, scalable storage across multiple Lambda functions and other AWS resources.
Lambda Mounting:EFS file systems can be mounted within Lambda functions to access a shared storage space.
Log Appending:EFS supports appending data to existing files, making it ideal for the log file scenario.
Amazon EFS Documentation:https://docs.aws.amazon.com/efs/
Using Amazon EFS with AWS Lambda:https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html

QUESTION 76
A company hosts its application on AWS. The application runs on an Amazon Elastic Container Service (Amazon ECS) cluster that uses AWS Fargate. The cluster runs behind an Application Load Balancer The application stores data in an Amazon Aurora database A developer encrypts and manages database credentials inside the application
The company wants to use a more secure credential storage method and implement periodic credential rotation.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate the secret credentials to Amazon RDS parameter groups. Encrypt the parameter by using an AWS Key Management Service (AWS KMS) key Turn on secret rotation. Use 1AM policies and roles to grant AWS KMS permissions to access Amazon RDS.
B. Migrate the credentials to AWS Systems Manager Parameter Store. Encrypt the parameter by using an AWS Key Management Service (AWS KMS) key. Turn on secret rotation. Use 1AM policies and roles to grant Amazon ECS Fargate permissions to access to AWS Secrets Manager
C. Migrate the credentials to ECS Fargate environment variables. Encrypt the credentials by using an AWS Key Management Service (AWS KMS) key Turn on secret rotation. Use 1AM policies and roles to grant Amazon ECS Fargate permissions to access to AWS Secrets Manager.
D. Migrate the credentials to AWS Secrets Manager. Encrypt the credentials by using an AWS Key Management Service (AWS KMS) key Turn on secret rotation Use 1AM policies and roles to grant Amazon ECS Fargate permissions to access to AWS Secrets Manager by using keys.

Correct Answer: D
Section:
Explanation:
Secrets Management:AWS Secrets Manager is designed specifically for storing and managing sensitive credentials.
Built-in Rotation:Secrets Manager provides automatic secret rotation functionality, enhancing security posture significantly.

IAM Integration:IAM policies and roles grant fine-grained access to ECS Fargate, ensuring the principle of least privilege.
Reduced Overhead:This solution centralizes secrets management and automates rotation, reducing operational overhead compared to the other options.
AWS Secrets Manager:https://aws.amazon.com/secrets-manager/
Secrets Manager Rotation:https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html
IAM for Secrets Manager:https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-and-access_iam-policies.html

QUESTION 77
A developer is testing a RESTful application that is deployed by using Amazon API Gateway and AWS Lambda When the developer tests the user login by using credentials that are not valid, the developer receives an HTTP 405 METHOD_NOT_ALLOWED error The developer has verified that the test is sending the correct request for the resource
Which HTTP error should the application return in response to the request?

A. HTTP 401
B. HTTP 404
C. HTTP 503
D. HTTP 505

Correct Answer: A
Section:
Explanation:
HTTP Status Codes:Each HTTP status code has a specific meaning in RESTful APIs.
HTTP 405 (Method Not Allowed):Indicates that the request method (e.g., POST) is not supported for the specified resource.
HTTP 401 (Unauthorized):Represents a failure to authenticate, which is the appropriate response for invalid login credentials.
HTTP Status Codes:https://developer.mozilla.org/en-US/docs/Web/HTTP/Status

QUESTION 78
A company runs an application on AWS The application uses an AWS Lambda function that is configured with an Amazon Simple Queue Service (Amazon SQS) queue called high priority queue as the event source A developer is updating the Lambda function with another SQS queue called low priority queue as the event source The Lambda function must always read up to 10 simultaneous messages from the high priority queue before processing messages from low priority queue. The Lambda function must be limited to 100 simultaneous invocations.
Which solution will meet these requirements'?

A. Set the event source mapping batch size to 10 for the high priority queue and to 90 for the low priority queue
B. Set the delivery delay to 0 seconds for the high priority queue and to 10 seconds for the low priority queue
C. Set the event source mapping maximum concurrency to 10 for the high priority queue and to 90 for the low priority queue
D. Set the event source mapping batch window to 10 for the high priority queue and to 90 for the low priority queue

Correct Answer: C
Section:
Explanation:
Lambda Concurrency:The 'maximum concurrency' setting in event source mappings controls the maximum number of simultaneous invocations Lambda allows for that specific source.
Prioritizing Queues:Setting a lower maximum concurrency for the 'high priority queue' ensures it's processed first while allowing more concurrent invocations from the 'low priority queue'.
Batching:Batch size settings affect the number of messages Lambda retrieves from a queue per invocation, which is less relevant to the prioritization requirement.
Lambda Event Source Mappings:https://docs.aws.amazon.com/lambda/latest/dg/invocation-eventsourcemapping.html
Lambda Concurrency:https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html

QUESTION 79
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?

A. Query the instance metadata from http://169.254.169.254. latestmeta-data/.
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

Correct Answer: A
Section:
Explanation:
Instance Metadata Service:EC2 instances have access to an internal metadata service. It provides instance-specific information like instance ID, security groups, and public IP address.
Accessing Metadata:
Make an HTTP GET request to the base URL:http://169.254.169.254/latest/meta-data/
You'll get a list of available categories. The public IPv4 address is underpublic-ipv4.
Instance Metadata and User Data:https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html

QUESTION 80
A company has a web application that is hosted on Amazon EC2 instances The EC2 instances are configured to stream logs to Amazon CloudWatch Logs The company needs to receive an Amazon Simple Notification Service (Amazon SNS) notification when the number of application error messages exceeds a defined threshold within a 5-minute period
Which solution will meet these requirements?

A. Rewrite the application code to stream application logs to Amazon SNS Configure an SNS topic to send a notification when the number of errors exceeds the defined threshold within a 5-minute period
B. Configure a subscription filter on the CloudWatch Logs log group. Configure the filter to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period.
C. Install and configure the Amazon Inspector agent on the EC2 instances to monitor for errors Configure Amazon Inspector to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period
D. Create a CloudWatch metric filter to match the application error pattern in the log data. Set up a CloudWatch alarm based on the new custom metric. Configure the alarm to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period.

Correct Answer: D
Section:
Explanation:
CloudWatch for Log Analysis:CloudWatch is the best fit here because logs are already centralized. Here's the process:
Metric Filter: Create a metric filter on the CloudWatch Logs log group. Design a pattern to specifically identify application error messages.
Custom Metric: This filter generates a new custom CloudWatch metric (e.g.,ApplicationErrors). This metric tracks the error count.
CloudWatch Alarm: Create an alarm on theApplicationErrorsmetric. Configure the alarm with your desired threshold and a 5-minute evaluation period.
SNS Action: Set the alarm to trigger an SNS notification when it enters the alarm state.
CloudWatch Metric Filters:https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html
CloudWatch Alarms:https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html

QUESTION 81
A developer is creating a service that uses an Amazon S3 bucket for image uploads. The service will use an AWS Lambda function to create a thumbnail of each image Each time an image is uploaded the service needs to send an email notification and create the thumbnail The developer needs to configure the image processing and email notifications setup.
Which solution will meet these requirements?

A. Create an Amazon Simple Notification Service (Amazon SNS) topic Configure S3 event notifications with a destination of the SNS topic Subscribe the Lambda function to the SNS topic Create an email notification subscription to the SNS topic
B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure S3 event notifications with a destination of the SNS topic. Subscribe the Lambda function to the SNS topic. Create an Amazon Simple Queue Service (Amazon SQS) queue Subscribe the SQS queue to the SNS topic Create an email notification subscription to the SQS queue.
C. Create an Amazon Simple Queue Service (Amazon SQS) queue Configure S3 event notifications with a destination of the SQS queue Subscribe the Lambda function to the SQS queue Create an email notification subscription to the SQS queue.

D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Send S3 event notifications to Amazon EventBridge. Create an EventBndge rule that runs the Lambda function when images are uploaded to the S3 bucket Create an EventBridge rule that sends notifications to the SQS queue Create an email notification subscription to the SQS queue

Correct Answer: A
Section:
Explanation:
SNS as a Fan-out Mechanism:SNS is perfect for triggering multiple actions from a single event (here, the image upload).
Workflow:
SNS Topic:Create an SNS topic that will be the central notification point.
S3 Event Notification:Configure the S3 bucket to send 'Object Created' event notifications to the SNS topic.
Lambda Subscription:Subscribe your thumbnail-creating Lambda function to the SNS topic.
Email Subscription:Subscribe an email address to the SNS topic to trigger notifications.
S3 Event Notifications:https://docs.aws.amazon.com/AmazonS3/latest/userguide/EventNotifications.html
SNS Subscriptions:https://docs.aws.amazon.com/sns/latest/dg/SNSMobilePush.html

QUESTION 82
A company is developing an application that will be accessed through the Amazon API Gateway REST API. Registered users should be the only ones who can access certain resources of this API. The token being used should expire automatically and needs to be refreshed periodically.
How can a developer meet these requirements?

A. Create an Amazon Cognito identity pool, configure the Amazon Cognito Authorizer in API Gateway, and use the temporary credentials generated by the identity pool.
B. Create and maintain a database record for each user with a corresponding token and use an AWS Lambda authorizer in API Gateway.
C. Create an Amazon Cognito user pool, configure the Cognito Authorizer in API Gateway, and use the identity or access token.
D. Create an 1AM user for each API user, attach an invoke permissions policy to the API. and use an I AM authorizer in API Gateway.

Correct Answer: C
Section:

QUESTION 83
A developer manages a website that distributes its content by using Amazon CloudFront. The website's static artifacts are stored in an Amazon S3 bucket.
The developer deploys some changes and can see the new artifacts in the S3 bucket. However, the changes do not appear on the webpage that the CloudFront distribution delivers.
How should the developer resolve this issue?

A. Configure S3 Object Lock to update to the latest version of the files every time an S3 object is updated.
B. Configure the S3 bucket to clear all old objects from the bucket before new artifacts are uploaded.
C. Set CloudFront to invalidate the cache after the artifacts have been deployed to Amazon S3.
D. Set CloudFront to modify the distribution origin after the artifacts have been deployed to Amazon S3.

Correct Answer: C
Section:

QUESTION 84
A company had an Amazon RDS for MySQL DB instance that was named mysql-db. The DB instance was deleted within the past 90 days. A developer needs to find which 1AM user or role deleted the DB instance in the AWS environment. Which solution will provide this information?

A. Retrieve the AWS CloudTrail events for the resource mysql-db where the event name is DeleteDBInstance. Inspect each event.
B. Retrieve the Amazon CloudWatch log events from the most recent log stream within the rds/mysql-db log group. Inspect the log events.
C. Retrieve the AWS X-Ray trace summaries. Filter by services with the name mysql-db. Inspect the ErrorRootCauses values within each summary.

D. Retrieve the AWS Systems Manager deletions inventory Filter the inventory by deletions that have a TypeName value of RDS. Inspect the deletion details.

Correct Answer: A
Section:

QUESTION 85
A company is using an Amazon API Gateway REST API endpoint as a webhook to publish events from an on-premises source control management (SCM) system to Amazon EventBridge. The company has configured an EventBridge rule to listen for the events and to control application deployment in a central AWS account. The company needs to receive the same events across multiple receiver AWS accounts.
How can a developer meet these requirements without changing the configuration of the SCM system?

A. Deploy the API Gateway REST API to all the required AWS accounts. Use the same custom domain name for all the gateway endpoints so that a single SCM webhook can be used for all events from all accounts.
B. Deploy the API Gateway REST API to all the receiver AWS accounts. Create as many SCM webhooks as the number of AWS accounts.
C. Grant permission to the central AWS account for EventBridge to access the receiver AWS accounts. Add an EventBridge event bus on the receiver AWS accounts as the targets to the existing EventBridge rule.
D. Convert the API Gateway type from REST API to HTTP API.

Correct Answer: C
Section:

QUESTION 86
A developer is creating AWS CloudFormation templates to manage an application's deployment in Amazon Elastic Container Service (Amazon ECS) through AWS CodeDeploy. The developer wants to automatically deploy new versions of the application to a percentage of users before the new version becomes available for all users.
How should the developer manage the deployment of the new version?

A. Modify the CloudFormation template to include a Transform section and the AWS::CodeDeploy::BlueGreen hook.
B. Deploy the new version in a new CloudFormation stack. After testing is complete, update the application's DNS records for the new stack.
C. Run CloudFormation stack updates on the application stack to deploy new application versions when they are available.
D. Create a nested stack for the new version. Include a Transform section and the AWS::CodeDeploy::BlueGreen hook.

Correct Answer: A
Section:

QUESTION 87
A developer is building a highly secure healthcare application using serverless components. This application requires writing temporary data to /Imp storage on an AWS Lambda function.
How should the developer encrypt this data?

A. Enable Amazon EBS volume encryption with an AWS KMS key in the Lambda function configuration so that all storage attached to the Lambda function is encrypted.
B. Set up the Lambda function with a role and key policy to access an AWS KMS key. Use the key to generate a data key used to encrypt all data prior to writing to Amp storage.
C. Use OpenSSL to generate a symmetric encryption key on Lambda startup. Use this key to encrypt the data prior to writing to /tmp.
D. Use an on-premises hardware security module (HSM) to generate keys, where the Lambda function requests a data key from the HSM and uses that to encrypt data on all requests to the function.

Correct Answer: B
Section:

QUESTION 88
A developer is creating an AWS Serverless Application Model (AWS SAM) template. The AWS SAM template contains the definition of multiple AWS Lambda functions, an Amazon S3 bucket, and an Amazon CtoudFront distribution. One of the Lambda functions runs on Lambda@Edge in the CloudFront distribution. The S3 bucket is configured as an origin for the CloudFront distribution.
When the developer deploys the AWS SAM template in the eu-west-1 Region, the creation of the stack fails.
Which of the following could be the reason for this issue?

A. CloudFront distributions can be created only in the us-east-1 Region.
B. Lambda@Edge functions can be created only in the us-east-1 Region.
C. A single AWS SAM template cannot contain multiple Lambda functions.
D. The CloudFront distribution and the S3 bucket cannot be created in the same Region.

Correct Answer: B
Section:

QUESTION 89
A developer has written a distributed application that uses micro services. The microservices are running on Amazon EC2 instances. Because of message volume, the developer is unable to match log output from each microservice to a specific transaction. The developer needs to analyze the message flow to debug the application.
Which combination of steps should the developer take to meet this requirement? (Select TWO.)

A. Download the AWS X-Ray daemon. Install the daemon on an EC2 instance. Ensure that the EC2 instance allows UDP traffic on port 2000.
B. Configure an interface VPC endpoint to allow traffic to reach the global AWS X-Ray daemon on TCP port 2000.
C. Enable AWS X-Ray. Configure Amazon CloudWatch to push logs to X-Ray.
D. Add the AWS X-Ray software development kit (SDK) to the microservices. Use X-Ray to trace requests that each microservice makes.
E. Set up Amazon CloudWatch metric streams to collect streaming data from the microservices.

Correct Answer: A, D
Section:

QUESTION 90
A developer is building an application that uses Amazon DynamoDB. The developer wants to retrieve multiple specific items from the database with a single API call. Which DynamoDB API call will meet these requirements with the MINIMUM impact on the database?

A. BatchGetItem
B. GetItem
C. Scan
D. Query

Correct Answer: A
Section:

QUESTION 91
An application stores user data in Amazon S3 buckets in multiple AWS Regions. A developer needs to implement a solution that analyzes the user data in the S3 buckets to find sensitive information. The analysis findings from all the S3 buckets must be available in the eu-west-2 Region.
Which solution will meet these requirements with the LEAST development effort?

A. Create an AWS Lambda function to generate findings. Program the Lambda function to send the findings to another S3 bucket in eu-west-2.
B. Configure Amazon Made to generate findings. Use Amazon EventBridge to create rules that copy the findings to eu-west-2.
C. Configure Amazon Inspector to generate findings. Use Amazon EventBridge to create rules that copy the findings to eu-west-2.
D. Configure Amazon Macie to generate findings and to publish the findings to AWS CloudTrail. Use a CloudTrail trail to copy the results to eu-west-2.

Correct Answer: B
Section:

QUESTION 92
A company uses Amazon DynamoDB as a data store for its order management system. The company frontend application stores orders in a DynamoDB table. The DynamoDB table is configured to send change events to a DynamoDB stream. The company uses an AWS Lambda function to log and process the incoming orders based on data from the DynamoDB stream.
An operational review reveals that the order quantity of incoming orders is sometimes set to 0. A developer needs to create a dashboard that will show how many unique customers this problem affects each day.
What should the developer do to implement the dashboard?

A. Grant the Lambda function's execution role permissions to upload logs to Amazon CloudWatch Logs. Implement a CloudWatch Logs Insights query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day periods. Add the CloudWatch Logs Insights query to a CloudWatch dashboard.
B. Use Amazon Athena to query AWS CtoudTrail API logs for API calls. Implement an Athena query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day periods. Add the Athena query to an Amazon CloudWatch dashboard.
C. Configure the Lambda function to send events to Amazon EventBridge. Create an EventBridge rule that groups the number of unique customers for orders with order quantity equal to 0 in 1-day periods. Add a CloudWatch dashboard as the target of the rule.
D. Turn on custom Amazon CloudWatch metrics for the DynamoDB stream of the DynamoOB table. Create a CloudWatch alarm that groups the number of unique customers for orders with order quantity equal to 0 in 1-day periods. Add the CloudWatch alarm to a CloudWatch dashboard.

Correct Answer: A
Section:

QUESTION 93
A developer is integrating Amazon ElastiCache in an application. The cache will store data from a database. The cached data must populate real-time dashboards. Which caching strategy will meet these requirements?

A. A read-through cache
B. A write-behind cache
C. A lazy-loading cache
D. A write-through cache

Correct Answer: D
Section:

QUESTION 94
A company's application has an AWS Lambda function that processes messages from IoT devices. The company wants to monitor the Lambda function to ensure that the Lambda function is meeting its required service level agreement (SLA).
A developer must implement a solution to determine the application's throughput in near real time. The throughput must be based on the number of messages that the Lambda function receives and processes in a given time period. The Lambda function performs initialization and post-processing steps that must not factor into the throughput measurement.
What should the developer do to meet these requirements?

A. Use the Lambda function's ConcurrentExecutions metric in Amazon CloudWatch to measure the throughput.
B. Modify the application to log the calculated throughput to Amazon CloudWatch Logs. Use Amazon EventBridge to invoke a separate Lambda function to process the logs on a schedule.
C. Modify the application to publish custom Amazon CloudWatch metrics when the Lambda function receives and processes each message. Use the metrics to calculate the throughput.
D. Use the Lambda function's Invocations metric and Duration metric to calculate the throughput in Amazon CloudWatch.

Correct Answer: C
Section:

QUESTION 95
A developer is using AWS CodeDeploy to launch an application onto Amazon EC2 instances. The application deployment fails during testing. The developer notices an IAM_ROLE_PERMISSIONS error code in Amazon CloudWatch logs.
What should the developer do to resolve the error?

A. Ensure that the deployment group is using the correct role name for the CodeDeploy service role.
B. Attach the AWSCodeDeployRoleECS policy to the CodeDeploy service role.
C. Attach the AWSCodeDeployRole policy to the CodeDeploy service role.
D. Ensure the CodeDeploy agent is installed and running on all instances in the deployment group.

Correct Answer: C
Section:

QUESTION 96
A company is building a serverless application that uses AWS Lambda functions. The company needs to create a set of test events to test Lambda functions in a development environment. The test events will be created once and then will be used by all the developers in an 1AM developer group. The test events must be editable by any of the 1AM users in the 1AM developer group.
Which solution will meet these requirements?

A. Create and store the test events in Amazon S3 as JSON objects. Allow S3 bucket access to all 1AM users.
B. Create the test events. Configure the event sharing settings to make the test events shareable.
C. Create and store the test events in Amazon DynamoDB. Allow access to DynamoDB by using 1AM roles.
D. Create the test events. Configure the event sharing settings to make the test events private.

Correct Answer: B
Section:

QUESTION 97
A developer has deployed an AWS Lambda function that is subscribed to an Amazon Simple Notification Service {Amazon SNS) topic. The developer must implement a solution to add a record of each Lambda function invocation to an Amazon Simple Queue Service {Amazon SQS) queue.
Which solution will meet this requirement?

A. Configure the SQS queue as a dead-letter queue for the Lambda function.
B. Create code that uses the AWS SDK to call the SQS SendMessage operation to add the invocation details to the SQS queue. Add the code to the end of the Lambda function.
C. Add two asynchronous invocation destinations to the Lambda function: one destination for successful invocations and one destination for failed invocations. Configure the SQS queue as the destination for each type. Create an Amazon CloudWatch alarm based on the DestinationDeliveryFailures metric to catch any message that cannot be delivered.
D. Add a single asynchronous invocation destination to the Lambda function to capture successful invocations. Configure the SQS queue as the destination. Create an Amazon CloudWatch alarm based on the DestinationDeliveryFailures metric to catch any message that cannot be delivered.

Correct Answer: D
Section:

QUESTION 98
A developer is creating an application that will be deployed on IoT devices. The application will send data to a RESTful API that is deployed as an AWS Lambda function. The application will assign each API request a unique identifier. The volume of API requests from the application can randomly increase at any given time of day.
During periods of request throttling, the application might need to retry requests. The API must be able to handle duplicate requests without inconsistencies or data loss.
Which solution will meet these requirements?

A. Create an Amazon RDS for MySQL DB instance. Store the unique identifier for each request in a database table. Modify the Lambda function to check the table for the identifier before processing the request.
B. Create an Amazon DynamoDB table. Store the unique identifier for each request in the table.
Modify the Lambda function to check the table for the identifier before processing the request.
C. Create an Amazon DynamoDB table. Store the unique identifier for each request in the table. Modify the Lambda function to return a client error response when the function receives a duplicate request.

D. Create an Amazon ElastiCache for Memcached instance. Store the unique identifier for each request in the cache. Modify the Lambda function to check the cache for the identifier before processing the request.

Correct Answer: B
Section:
Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that can store and retrieve any amount of data with high availability and performance. DynamoDB can handle concurrent requests from multiple IoT devices without throttling or data loss. To prevent duplicate requests from causing inconsistencies or data loss, the Lambda function can use DynamoDB conditional writes to check if the unique identifier for each request already exists in the table before processing the request. If the identifier exists, the function can skip or abort the request; otherwise, it can process the request and store the identifier in the table. Reference: Using conditional writes

QUESTION 99
A developer wants to expand an application to run in multiple AWS Regions. The developer wants to copy Amazon Machine Images (AMIs) with the latest changes and create a new application stack in the destination Region. According to company requirements, all AMIs must be encrypted in all Regions. However, not all the AMIs that the company uses are encrypted.
How can the developer expand the application to run in the destination Region while meeting the encryption requirement?

A. Create new AMIs, and specify encryption parameters. Copy the encrypted AMIs to the destination Region. Delete the unencrypted AMIs.
B. Use AWS Key Management Service (AWS KMS) to enable encryption on the unencrypted AMIs. Copy the encrypted AMIs to the destination Region.
C. Use AWS Certificate Manager (ACM) to enable encryption on the unencrypted AMIs. Copy the encrypted AMIs to the destination Region.
D. Copy the unencrypted AMIs to the destination Region. Enable encryption by default in the destination Region.

Correct Answer: A
Section:
Explanation:
Amazon Machine Images (AMIs) are encrypted snapshots of EC2 instances that can be used to launch new instances. The developer can create new AMIs from the existing instances and specify encryption parameters. The developer can copy the encrypted AMIs to the destination Region and use them to create a new application stack. The developer can delete the unencrypted AMIs after the encryption process is complete. This solution will meet the encryption requirement and allow the developer to expand the application to run in the destination Region.
Reference:
[Amazon Machine Images (AMI) - Amazon Elastic Compute Cloud]
[Encrypting an Amazon EBS Snapshot - Amazon Elastic Compute Cloud]
[Copying an AMI - Amazon Elastic Compute Cloud]

QUESTION 100
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 bucket. Assign an access point to each web application bucket.
B. Create a bucket policy that allows access to the central S3 bucket. Attach the bucket policy to the central S3 bucket.
C. Create a cross-origin resource sharing (CORS) configuration that allows access to the central S3 bucket. Add the CORS configuration to the central S3 bucket.
D. Create a Content-MD5 header that provides a message integrity check for the central S3 bucket.
Insert the Content-MD5 header for each web application request.

Correct Answer: C
Section:
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

QUESTION 101
An application is processing clickstream data using Amazon Kinesis. The clickstream data feed into Kinesis experiences periodic spikes. The PutRecords API call occasionally fails and the logs show that the failed call returns the response shown below:

```
{
   "FailedRecordCount": 1,
   "Records": [
       {
           "SequenceNumber": "21269319989900637946712965403778482371",
           "ShardId": "shardId-000000000001"
       },
       {
           "ErrorCode":"ProvisionedThroughputExceededException",
           "ErrorMessage": "Rate exceeded for shard shardId-000000000001 in

                stream exampleStreamName under account 123456789."
       },
       {
           "SequenceNumber": "21269319989999637946712965403778482985",
           "ShardId": "shardId-000000000002"
       }
   ]
}
```

Which techniques will help mitigate this exception? (Choose two.)

A. Implement retries with exponential backoff.
B. Use a PutRecord API instead of PutRecords.
C. Reduce the frequency and/or size of the requests.
D. Use Amazon SNS instead of Kinesis.
E. Reduce the number of KCL consumers.

Correct Answer: A, C
Section:
Explanation:
The response from the API call indicates that the ProvisionedThroughputExceededException exception has occurred. This exception means that the rate of incoming requests exceeds the throughput limit for one or more shards in a stream. To mitigate this exception, the developer can use one or more of the following techniques:
Implement retries with exponential backoff. This will introduce randomness in the retry intervals and avoid overwhelming the shards with retries.
Reduce the frequency and/or size of the requests. This will reduce the load on the shards and avoid throttling errors.
Increase the number of shards in the stream. This will increase the throughput capacity of the stream and accommodate higher request rates.
Use a PutRecord API instead of PutRecords. This will reduce the number of records per request and avoid exceeding the payload limit.
Reference:
[ProvisionedThroughputExceededException - Amazon Kinesis Data Streams Service API Reference]
[Best Practices for Handling Kinesis Data Streams Errors]

QUESTION 102
A company has an Amazon S3 bucket that contains sensitive dat a. The data must be encrypted in transit and at rest. The company encrypts the data in the S3 bucket by using an AWS Key Management Service (AWS KMS) key.
A developer needs to grant several other AWS accounts the permission to use the S3 GetObject operation to retrieve the data from the S3 bucket.
How can the developer enforce that all requests to retrieve the data provide encryption in transit?

A. Define a resource-based policy on the S3 bucket to deny access when a request meets the condition "aws:SecureTransport": "false".
B. Define a resource-based policy on the S3 bucket to allow access when a request meets the condition "aws:SecureTransport": "false".
C. Define a role-based policy on the other accounts' roles to deny access when a request meets the condition of "aws:SecureTransport": "false".
D. Define a resource-based policy on the KMS key to deny access when a request meets the condition of "aws:SecureTransport": "false".

Correct Answer: A
Section:
Explanation:
Amazon S3 supports resource-based policies, which are JSON documents that specify the permissions for accessing S3 resources. A resource-based policy can be used to enforce encryption in transit by denying access to requests that do not use HTTPS. The condition key aws:SecureTransport can be used to check if the request was sent using SSL. If the value of this key is false, the request is denied; otherwise, the request is allowed.
Reference: How do I use an S3 bucket policy to require requests to use Secure Socket Layer (SSL)?

QUESTION 103
A development learn has an Amazon API Gateway REST API that is backed by an AWS Lambda function.
Users have reported performance issues for the Lambda function. The development team identified the source of the issues as a cold start of the Lambda function. The development team needs to reduce the time needed for the Lambda function to initialize.
Which solution will meet this requirement?

A. Change the Lambda concurrency lo reserved concurrency.
B. Increase the timeout of the Lambda function.
C. Increase the memory allocation of the Lambda function.
D. Configure provisioned concurrency for the Lambda function.

Correct Answer: D
Section:

QUESTION 104
A developer is creating a stock trading application. The developer needs a solution to send text messages to application users to confirmation when a trade has been completed.
The solution must deliver messages in the order a user makes stock trades. The solution must not send duplicate messages.
Which solution will meet these requirements?

A. Configure the application to publish messages to an Amazon Data Firehose delivery stream. Configure the delivery stream to have a destination of each user's mobile phone number that is passed in the trade confirmation message.
B. Create an Amazon Simple Queue Service (Amazon SQS) FIFO queue. Use the SendMessageln API call to send the trade confirmation messages to the queue. Use the SendMessageOut API to send the messages to users by using the information provided in the trade confirmation message.
C. Configure a pipe in Amazon EventBridge Pipes. Connect the application to the pipe as a source. Configure the pipe to use each user's mobile phone number as a target. Configure the pipe to send incoming events to the users.
D. Create an Amazon Simple Notification Service (SNS) FIFO topic. Configure the application to use the AWS SDK to publish notifications to the SNS topic to send SMS messages to the users.

Correct Answer: C
Section:

QUESTION 105
A company offers a business-to-business software service that runs on dedicated infrastructure deployed in each customer's AWS account. Before a feature release, the company needs to run integration tests on real AWS test infrastructure. The test infrastructure consists of Amazon EC2 instances and an Amazon RDS database.
A developer must set up a continuous delivery process that will provision the test infrastructure across the different AWS accounts. The developer then must run the integration tests.
Which solution will meet these requirements with the LEAST administrative effort?

A. Use AWS CodeDeploy with AWS CloudFormation StackSets to deploy the infrastructure. Use Amazon CodeGuru to run the tests.
B. Use AWS CodePipeline with AWS CloudFormation StackSets to deploy the infrastructure. Use AWS CodeBuild to run the tests.
C. Use AWS CodePipeline with AWS CloudFormation change sets to deploy the infrastructure. Use a CloudFormation custom resource to run the tests.
D. Use AWS Serverless Application Model (AWS SAM) templates with AWS CloudFormation change sets to deploy the infrastructure. Use AWS CodeDeploy to run the tests.

Correct Answer: B
Section:

QUESTION 106
A developer is making changes to a custom application that uses AWS Elastic Beanstalk.
Which solutions will update the Elastic Beanstalk environment with the new application version after the developer completes the changes? (Select TWO.)

A. Package the application code into a .zip file. Use the AWS Management Console to upload the .zip file and deploy the packaged application.
B. Package the application code into a .tar file. Use the AWS Management Console to create a new application version from the .tar file. Update the environment by using the AWS CLI.
C. Package the application code into a .tar file. Use the AWS Management Console to upload the .tar file and deploy the packaged application.
D. Package the application code into a .zip file. Use the AWS CLI to create a new application version from the .zip file and to update the environment.
E. Package the application code into a .zip file. Use the AWS Management Console to create a new application version from the .zip file. Rebuild the environment by using the AWS CLI.

Correct Answer: A, D
Section:

QUESTION 107
A company has an AWS Step Functions state machine named myStateMachine. The company configured a service role for Step Functions. The developer must ensure that only the myStateMachine state machine can assume the service role.

A. 'Condition': { 'ArnLike': { 'aws ':'arn:aws:states:ap-south-1:111111111111:stateMachine ' } }
B. ' Condition': { 'ArnLike': { 'aws ':'arn:aws:states:ap-south-1:*:stateMachine ' } }

Correct Answer: A
Section:
Explanation:
Comprehensive Detailed Step by Step Explanation with All AWS Developer
Reference: To ensure that only a specific AWS Step Functions state machine (myStateMachine) can assume the service role, you must configure the correct trust policy in AWS IAM.
Trust Policies: Trust policies determine which entities (services or users) are allowed to assume the role. In this case, we want to restrict the trust policy to only allow the specific state machine (myStateMachine) to assume the role.
Using ArnLike: The condition 'ArnLike' is used to specify that the SourceArn (which refers to the ARN of the entity assuming the role) must match a specific ARN. Option A specifies the exact ARN of the myStateMachine state machine, ensuring that only this state machine can assume the role.
Option B: This option is incorrect because it uses a wildcard (*) for the account ID, which would allow any state machine in the ap-south-1 region to assume the role, not just the specific one.
AWS Step Functions IAM Policies

QUESTION 108
A developer used the AWS SDK to create an application that aggregates and produces log records for 10 services. The application delivers data to an Amazon Kinesis Data Streams stream.
Each record contains a log message with a service name, creation timestamp, and other log information. The stream has 15 shards in provisioned capacity mode. The stream uses service name as the partition key.
The developer notices that when all the services are producing logs, ProvisionedThroughputExceededException errors occur during PutRecord requests. The stream metrics show that the write capacity the applications use is below the provisioned capacity.

A. Change the capacity mode from provisioned to on-demand.

B. Double the number of shards until the throttling errors stop occurring.
C. Change the partition key from service name to creation timestamp.
D. Use a separate Kinesis stream for each service to generate the logs.

Correct Answer: C
Section:
Explanation:
Comprehensive and Detailed Step-by-Step
Issue Analysis:
The stream uses service name as the partition key. This can cause 'hot partition' issues when a few service names generate significantly more logs compared to others, causing uneven distribution of data across shards.
Metrics show that the write capacity used is below provisioned capacity, which confirms that the throughput errors are due to shard-level limits and not overall capacity.
Option C: Change Partition Key to Creation Timestamp:
By changing the partition key to the creation timestamp (or a composite key including timestamp), the distribution of data across shards can be randomized, ensuring an even spread of records.
This resolves the shard overutilization issue and eliminates ProvisionedThroughputExceededException.
Why Other Options Are Incorrect:
Option A: Switching to on-demand capacity mode might temporarily alleviate the issue, but the root cause (hot partitioning) remains unresolved.
Option B: Adding shards increases capacity but does not fix the skewed data distribution caused by using the service name as the partition key.
Option D: Creating separate streams for each service adds unnecessary complexity and does not scale well as the number of services grows.
Best Practices for Kinesis Data Streams Partition Key Design

QUESTION 109
A development team is creating a serverless application that uses AWS Lambda functions. The team wants to streamline a testing workflow by sharing test events across multiple developers within the same AWS account. The team wants to ensure all developers can use consistent test events without compromising security.

A. Export test events as JSON files. Store the files in an Amazon S3 bucket. Configure granular IAM permissions to allow the developers to access the S3 bucket.
B. Store test events in an Amazon DynamoDB table. Create an AWS Lambda function to retrieve shared test events for the developers.
C. Configure test events to be shareable. Configure granular IAM permissions to allow the developers to access shared test events.
D. Set up a Git repository to store test events. Provide the developers with access to the repository.

Correct Answer: A
Section:
Explanation:
Comprehensive and Detailed Step-by-Step
Option A: Use Amazon S3 for Shared Test Events:
Storing JSON test event files in an S3 bucket provides a centralized, cost-effective, and highly available solution.
Granular IAM policies can restrict access to specific developers or roles, ensuring security while maintaining consistency for shared test events.
This solution has minimal operational overhead and integrates easily with existing workflows.
Why Other Options Are Incorrect:
Option B: Using DynamoDB and a Lambda function introduces unnecessary complexity for a relatively simple requirement. S3 provides a simpler and more cost-efficient solution.
Option C: AWS Lambda test events are not inherently shareable across developers, making this option invalid.
Option D: Using a Git repository adds operational overhead and requires developers to clone/update repositories for access, which is more cumbersome compared to S3.
Amazon S3 for Centralized Storage

QUESTION 110
A developer needs to export the contents of several Amazon DynamoDB tables into Amazon S3 buckets to comply with company data regulations. The developer uses the AWS CLI to run commands to export from each table to the proper S3 bucket. The developer sets up AWS credentials correctly and grants resources appropriate permissions. However, the exports of some tables fail.
What should the developer do to resolve this issue?

A. Ensure that point-in-time recovery is enabled on the DynamoDB tables.

B. Ensure that the target S3 bucket is in the same AWS Region as the DynamoDB table.
C. Ensure that DynamoDB streaming is enabled for the tables.
D. Ensure that DynamoDB Accelerator (DAX) is enabled.

Correct Answer: B
Section:
Explanation:
Comprehensive Detailed and Lengthy Step-by-Step Explanation with All AWS Developer
Reference:
1. Understanding the Use Case:
The developer needs to export DynamoDB table data into Amazon S3 buckets using the AWS CLI, and some exports are failing. Proper credentials and permissions have already been configured.
2. Key Conditions to Check:
Region Consistency:
DynamoDB exports require that the target S3 bucket and the DynamoDB table reside in the same AWS Region. If they are not in the same Region, the export process will fail.
Point-in-Time Recovery (PITR):
PITR is not required for exporting data from DynamoDB to S3. Enabling PITR allows recovery of table states at specific points in time but does not directly influence export functionality.
DynamoDB Streams:
Streams allow real-time capture of data modifications but are unrelated to the bulk export feature.
DAX (DynamoDB Accelerator):
DAX is a caching service that speeds up read operations for DynamoDB but does not affect the export functionality.
3. Explanation of the Options:
Option A:
'Ensure that point-in-time recovery is enabled on the DynamoDB tables.'
While PITR is useful for disaster recovery and restoring table states, it is not required for exporting data to S3. This option does not address the export failure.
Option B:
'Ensure that the target S3 bucket is in the same AWS Region as the DynamoDB table.'
This is the correct answer. DynamoDB export functionality requires the target S3 bucket to reside in the same AWS Region as the DynamoDB table. If the S3 bucket is in a different Region, the export will fail.
Option C:
'Ensure that DynamoDB streaming is enabled for the tables.'
Streams are useful for capturing real-time changes in DynamoDB tables but are unrelated to the export functionality. This option does not resolve the issue.
Option D:
'Ensure that DynamoDB Accelerator (DAX) is enabled.'
DAX accelerates read operations but does not influence the export functionality. This option is irrelevant to the issue.
4. Resolution Steps:
To ensure successful exports:
Verify the Region of the DynamoDB tables:
Check the Region where each table is located.
Verify the Region of the target S3 buckets:
Confirm that the target S3 bucket for each export is in the same Region as the corresponding DynamoDB table.
If necessary, create new S3 buckets in the appropriate Regions.
Run the export command again with the correct setup:
aws dynamodb export-table-to-point-in-time \
--table-name <TableName> \
--s3-bucket <BucketName> \
--s3-prefix <Prefix> \
--export-time <ExportTime> \
--region <Region>
Exporting DynamoDB Data to Amazon S3
S3 Bucket Region Requirements for DynamoDB Exports
AWS CLI Reference for DynamoDB Export

QUESTION 111
A developer is creating an application that must be able to generate API responses without backend integrations. Multiple internal teams need to work with the API while the application is still in development. Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon API Gateway REST API. Set up a proxy resource that has the HTTP proxy integration type.
B. Create an Amazon API Gateway HTTP API. Provision a VPC link, and set up a private integration on the API to connect to a VPC.
C. Create an Amazon API Gateway HTTP API. Enable mock integration on the method of the API resource.
D. Create an Amazon API Gateway REST API. Enable mock integration on the method of the API resource.

Correct Answer: D
Section:
Explanation:
Comprehensive Detailed and Lengthy Step-by-Step Explanation with All AWS Developer
Reference:
1. Understanding the Use Case:
The API needs to:
Generate responses without backend integrations: This indicates the use of mock responses for testing.
Be used by multiple internal teams during development.
Minimize operational overhead.
2. Key Features of Amazon API Gateway:
REST APIs: Fully managed API Gateway option that supports advanced capabilities like mock integrations, request/response transformation, and more.
HTTP APIs: Lightweight option for building APIs quickly. It supports fewer features but has lower operational complexity and cost.
Mock Integration: Allows API Gateway to return pre-defined responses without requiring backend integration.
3. Explanation of the Options:
Option A:
'Create an Amazon API Gateway REST API. Set up a proxy resource that has the HTTP proxy integration type.'
A proxy integration requires a backend service for handling requests. This does not meet the requirement of 'no backend integrations.'
Option B:
'Create an Amazon API Gateway HTTP API. Provision a VPC link, and set up a private integration on the API to connect to a VPC.'
This requires setting up a VPC and provisioning resources, which increases operational overhead and is unnecessary for this use case.
Option C:
'Create an Amazon API Gateway HTTP API. Enable mock integration on the method of the API resource.'
While HTTP APIs can enable mock integrations, they have limited support for advanced features compared to REST APIs, such as detailed request/response customization. REST APIs are better suited for development environments requiring mock responses.
Option D:
'Create an Amazon API Gateway REST API. Enable mock integration on the method of the API resource.'
This is the correct answer. REST APIs with mock integration allow defining pre-configured responses directly within API Gateway, making them ideal for scenarios where backend services are unavailable. It provides flexibility for testing while minimizing operational overhead.
4. Implementation Steps:
To enable mock integration with REST API:
Create a REST API in API Gateway:
Open the API Gateway Console.
Choose Create API > REST API.
Define the API Resource and Methods:
Add a resource and method (e.g., GET or POST).
Set Up Mock Integration:
Select the method, and in the Integration Type, choose Mock Integration.
Configure the Mock Response:
Define a 200 OK response with the desired response body and headers.

Deploy the API:
Deploy the API to a stage (e.g., dev) to make it accessible.
5. Why REST API Over HTTP API?
REST APIs support detailed request/response transformations and robust mock integration features, which are ideal for development and testing scenarios.
While HTTP APIs offer lower cost and simplicity, they lack some advanced features required for fine-tuned mock integrations.
Amazon API Gateway REST API Features
Mock Integration in API Gateway
Comparison of REST and HTTP APIs in API Gateway