Number: DVA-C02
Passing Score: 800
Time Limit: 120
File Version: 23.0

Exam Code: DVA-C02

Exam Name: AWS Certified Developer - Associate

Exam A

QUESTION 1
A development team maintains a web application by using a single AWS CloudFormation template.
The template defines web servers and an Amazon RDS database. The team uses the Cloud Formation template to deploy the Cloud Formation stack to different environments.
During a recent application deployment, a developer caused the primary development database to be dropped and recreated. The result of this incident was a loss of dat a. The team needs to avoid accidental database deletion in the future.
Which solutions will meet these requirements? (Choose two.)

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

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

QUESTION 2
A company is implementing an application on Amazon EC2 instances. The application needs to process incoming transactions. When the application detects a transaction that is not valid, the application must send a chat message to the company's support team. To send the message, the application needs to retrieve the access token to authenticate by using the chat API.
A developer needs to implement a solution to store the access token. The access token must be encrypted at rest and in transit. The access token must also be accessible from other AWS accounts.
Which solution will meet these requirements with the LEAST management overhead?

A. Use an AWS Systems Manager Parameter Store SecureString parameter that uses an AWS Key Management Service (AWS KMS) AWS managed key to store the access token. Add a resource-based policy to the parameter to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Parameter Store. Retrieve the token from Parameter Store with the decrypt flag enabled. Use the decrypted access token to send the message to the chat.
B. Encrypt the access token by using an AWS Key Management Service (AWS KMS) customer managed key. Store the access token in an Amazon DynamoDB table. Update the IAM role of the EC2 instances with permissions to access DynamoDB and AWS KMS. Retrieve the token from DynamoDB.
Decrypt the token by using AWS KMS on the EC2 instances. Use the decrypted access token to send the message to the chat.
C. Use AWS Secrets Manager with an AWS Key Management Service (AWS KMS) customer managed key to store the access token. Add a resource-based policy to the secret to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Secrets Manager.
Retrieve the token from Secrets Manager. Use the decrypted access token to send the message to the chat.
D. Encrypt the access token by using an AWS Key Management Service (AWS KMS) AWS managed key. Store the access token in an Amazon S3 bucket. Add a bucket policy to the S3 bucket to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Amazon S3 and AWS KMS. Retrieve the token from the S3 bucket. Decrypt the token by using AWS KMS on the EC2 instances. Use the decrypted access token to send the massage to the chat.

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/secrets-manager-share-betweenaccounts/
https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-andaccess_ examples_cross.html

QUESTION 3
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 4
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 5
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 6
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 7
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 8
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 9
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 10
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 11

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 12
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 13
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 14
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 15
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 16
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 17
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 18
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 19
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 20
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 21
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 22
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 23
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 24
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 25
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 26
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 27
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 28
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 29
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 30
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 31
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 32
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 33
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 34
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 35
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 36
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 37
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 38
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."
       },
       {
           "SeguenceNumber": "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 39
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 40
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 41
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 42
A company notices that credentials that the company uses to connect to an external software as a service (SaaS) vendor are stored in a configuration file as plaintext.
The developer needs to secure the API credentials and enforce automatic credentials rotation on a quarterly basis.

Which solution will meet these requirements MOST securely?

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

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

QUESTION 43
An application that runs on AWS receives messages from an Amazon Simple Queue Service (Amazon SQS) queue and processes the messages in batches. The application sends the data to another SQS queue to be consumed by another legacy application. The legacy system can take up to 5 minutes to process some transaction data.
A developer wants to ensure that there are no out-of-order updates in the legacy system. The developer cannot alter the behavior of the legacy system.
Which solution will meet these requirements?

A. Use an SQS FIFO queue. Configure the visibility timeout value.
B. Use an SQS standard queue with a SendMessageBatchRequestEntry data type. Configure the DelaySeconds values.
C. Use an SQS standard queue with a SendMessageBatchRequestEntry data type. Configure the visibility timeout value.
D. Use an SQS FIFO queue. Configure the DelaySeconds value.

Correct Answer: A
Section:
Explanation:
An SQS FIFO queue is a type of queue that preserves the order of messages and ensures that each message is delivered and processed only once1. This is suitable for the scenario where the developer wants to ensure that there are no out-of-order updates in the legacy system.
The visibility timeout value is the amount of time that a message is invisible in the queue after a consumer receives it2. This prevents other consumers from processing the same message simultaneously.If the consumer does not delete the message before the visibility timeout expires, the message becomes visible again and another consumer can receive it2.
In this scenario, the developer needs to configure the visibility timeout value to be longer than the maximum processing time of the legacy system, which is 5 minutes. This will ensure that the message remains invisible in the queue until the legacy system finishes processing it and deletes it. This will prevent duplicate or out-of-order processing of messages by the legacy system.

QUESTION 44
A developer is troubleshooting an application in an integration environment. In the application, an Amazon Simple Queue Service (Amazon SQS) queue consumes messages and then an AWS Lambda function processes the messages. The Lambda function transforms the messages and makes an API call to a third-party service.
There has been an increase in application usage. The third-party API frequently returns an HTTP 429 Too Many Requests error message. The error message prevents a significant number of messages from being processed successfully.
How can the developer resolve this issue?

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

D. Configure maximum concurrency on the SQS event source based on the third-party service's documented rate limits.

Correct Answer: D
Section:
Explanation:
Maximum concurrency for SQS as an event source allows customers to control the maximum concurrent invokes by the SQS event source1.When multiple SQS event sources are configured to a function, customers can control the maximum concurrent invokes of individual SQS event source1.
In this scenario, the developer needs to resolve the issue of the third-party API frequently returning an HTTP 429 Too Many Requests error message, which prevents a significant number of messages from being processed successfully. To achieve this, the developer can follow these steps:
Find out the documented rate limits of the third-party API, which specify how many requests can be made in a given time period.
Configure maximum concurrency on the SQS event source based on the rate limits of the third-party API. This will limit the number of concurrent invokes by the SQS event source and prevent exceeding the rate limits of the third-party API.
Test and monitor the application performance and adjust the maximum concurrency value as needed.
By using this solution, the developer can reduce the frequency of HTTP 429 errors and improve the message processing success rate. The developer can also avoid throttling or blocking by the third-party API.

QUESTION 45
An online sales company is developing a serverless application that runs on AWS. The application uses an AWS Lambda function that calculates order success rates and stores the data in an Amazon DynamoDB table. A developer wants an efficient way to invoke the Lambda function every 15 minutes.
Which solution will meet this requirement with the LEAST development effort?

A. Create an Amazon EventBridge rule that has a rate expression that will run the rule every 15 minutes. Add the Lambda function as the target of the EventBridge rule.
B. Create an AWS Systems Manager document that has a script that will invoke the Lambda function on Amazon EC2. Use a Systems Manager Run Command task to run the shell script every 15 minutes.
C. Create an AWS Step Functions state machine. Configure the state machine to invoke the Lambda function execution role at a specified interval by using a Wait state. Set the interval to 15 minutes.
D. Provision a small Amazon EC2 instance. Set up a cron job that invokes the Lambda function every 15 minutes.

Correct Answer: A
Section:
Explanation:
The best solution for this requirement is option
A) Creating an Amazon EventBridge rule that has a rate expression that will run the rule every 15 minutes and adding the Lambda function as the target of the EventBridge rule is the most efficient way to invoke the Lambda function periodically.This solution does not require any additional resources or development effort, and it leverages the built-in scheduling capabilities of EventBridge1.

QUESTION 46
A developer is migrating an application to Amazon Elastic Kubernetes Service (Amazon EKS). The developer migrates the application to Amazon Elastic Container Registry (Amazon ECR) with an EKS cluster.
As part of the application migration to a new backend, the developer creates a new AWS account. The developer makes configuration changes to the application to point the application to the new AWS account and to use new backend resources. The developer successfully tests the changes within the application by deploying the pipeline.
The Docker image build and the pipeline deployment are successful, but the application is still connecting to the old backend. The developer finds that the application's configuration is still referencing the original EKS cluster and not referencing the new backend resources.
Which reason can explain why the application is not connecting to the new resources?

A. The developer did not successfully create the new AWS account.
B. The developer added a new tag to the Docker image.
C. The developer did not update the Docker image tag to a new version.
D. The developer pushed the changes to a new Docker image tag.

Correct Answer: C
Section:
Explanation:
The correct answer is C) The developer did not update the Docker image tag to a new version.

C) The developer did not update the Docker image tag to a new version. This is correct. When deploying an application to Amazon EKS, the developer needs to specify the Docker image tag that contains the application code and configuration. If the developer does not update the Docker image tag to a new version after making changes to the application, the EKS cluster will continue to use the old Docker image tag that references the original backend resources. To fix this issue, the developer should update the Docker image tag to a new version and redeploy the application to the EKS cluster.
A) The developer did not successfully create the new AWS account. This is incorrect. The creation of a new AWS account is not related to the application's connection to the backend resources. The developer can use any AWS account to host the EKS cluster and the backend resources, as long as they have the proper permissions and configurations.
B) The developer added a new tag to the Docker image. This is incorrect. Adding a new tag to the Docker image is not enough to deploy the changes to the application. The developer also needs to update the Docker image tag in the EKS cluster configuration, so that the EKS cluster can pull and run the new Docker image.
D) The developer pushed the changes to a new Docker image tag. This is incorrect. Pushing the changes to a new Docker image tag is not enough to deploy the changes to the application. The developer also needs to update the Docker image tag in the EKS cluster configuration, so that the EKS cluster can pull and run the new Docker image.
1: Amazon EKS User Guide, ''Deploying applications to your Amazon EKS cluster'', https://docs.aws.amazon.com/eks/latest/userguide/deploying-applications.html
2: Amazon ECR User Guide, ''Pushing an image'', https://docs.aws.amazon.com/AmazonECR/latest/userguide/docker-push-ecr-image.html
3: Amazon EKS User Guide, ''Updating an Amazon EKS cluster'', https://docs.aws.amazon.com/eks/latest/userguide/update-cluster.html

QUESTION 47
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment, the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

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

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

QUESTION 48
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions.
When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

A. Set the Deployment Preference Type to Canary10Percent10Minutes. Set the AutoPublishAlias property to the Lambda alias.
B. Set the Deployment Preference Type to LinearlOPercentEvery10Minutes. Set AutoPubIishAlias property to the Lambda alias.
C. Set the Deployment Preference Type to CanarylOPercentIOMinutes. Set the PreTraffic and PostTraffic properties to the Lambda alias.
D. Set the Deployment Preference Type to LinearlOPercentEverylOMinutes. Set PreTraffic and Post Traffic properties to the Lambda alias.

Correct Answer: A

Section:
Explanation:
The AWS Serverless Application Model (AWS SAM) comes built-in with CodeDeploy to provide gradual AWS Lambda deployments1. TheDeploymentPreferenceproperty in AWS SAM allows you to specify the type of deployment that you want. TheCanary10Percent10Minutesoption means that 10 percent of your customer traffic is immediately shifted to your new version.After 10 minutes, all traffic is shifted to the new version1.TheAutoPublishAliasproperty in AWS SAM allows AWS SAM to automatically create an alias that points to the updated version of the Lambda function1. Therefore, option A is correct.

QUESTION 49
A company developed an API application on AWS by using Amazon CloudFront, Amazon API Gateway, and AWS Lambd a. The API has a minimum of four requests every second. A developer notices that many API users run the same query by using the POST method. The developer wants to cache the POST request to optimize the API resources.
Which solution will meet these requirements?

A. Configure the CloudFront cache. Update the application to return cached content based upon the default request headers.
B. Override the cache method in the selected stage of API Gateway. Select the POST method.
C. Save the latest request response in Lambda /tmp directory. Update the Lambda function to check the /tmp directory.
D. Save the latest request in AWS Systems Manager Parameter Store. Modify the Lambda function to take the latest request response from Parameter Store.

Correct Answer: B
Section:
Explanation:
Amazon API Gateway provides tools for creating and documenting web APIs that route HTTP requests to Lambda functions2. You can secure access to your API with authentication and authorization controls.Your APIs can serve traffic over the internet or can be accessible only within your VPC2.You can override the cache method in the selected stage of API Gateway2. Therefore, option B is correct.

QUESTION 50
A company is building a compute-intensive application that will run on a fleet of Amazon EC2 instances. The application uses attached Amazon Elastic Block Store (Amazon EBS) volumes for storing data. The Amazon EBS volumes will be created at time of initial deployment. The application will process sensitive information. All of the data must be encrypted. The solution should not impact the application's performance.
Which solution will meet these requirements?

A. Configure the fleet of EC2 instances to use encrypted EBS volumes to store data.
B. Configure the application to write all data to an encrypted Amazon S3 bucket.
C. Configure a custom encryption algorithm for the application that will encrypt and decrypt all data.
D. Configure an Amazon Machine Image (AMI) that has an encrypted root volume and store the data to ephemeral disks.

Correct Answer: A
Section:
Explanation:
Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with Amazon EC2 instances1.Amazon EBS encryption offers a straight-forward encryption solution for your EBS resources associated with your EC2 instances1.When you create an encrypted EBS volume and attach it to a supported instance type, the following types of data are encrypted: Data at rest inside the volume, all data moving between the volume and the instance, all snapshots created from the volume, and all volumes created from those snapshots1. Therefore, option A is correct.

QUESTION 51
A developer is creating a new REST API by using Amazon API Gateway and AWS Lambd
a. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.
The developer wants to make the REST API available for testing by using API Gateway locally.
Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

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

D. Sam local start-api

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

QUESTION 52
A developer is creating an AWS Lambda function that consumes messages from an Amazon Simple Queue Service (Amazon SQS) standard queue. The developer notices that the Lambda function processes some messages multiple times.
How should developer resolve this issue MOST cost-effectively?

A. Change the Amazon SQS standard queue to an Amazon SQS FIFO queue by using the Amazon SQS message deduplication ID.
B. Set up a dead-letter queue.
C. Set the maximum concurrency limit of the AWS Lambda function to 1
D. Change the message processing to use Amazon Kinesis Data Streams instead of Amazon SQS.

Correct Answer: A
Section:
Explanation:
Amazon Simple Queue Service (Amazon SQS) is a fully managed queue service that allows you to de-couple and scale for applications1.Amazon SQS offers two types of queues: Standard and FIFO (First In First Out) queues1.The FIFO queue uses themessageDeduplicationIdproperty to treat messages with the same value as duplicate2. Therefore, changing the Amazon SQS standard queue to an Amazon SQS FIFO queue using the Amazon SQS message deduplication ID can help resolve the issue of the Lambda function processing some messages multiple times. Therefore, option A is correct.

QUESTION 53
A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node.js application.
To minimize these bugs, the developer wants to implement 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 (CI/CD) pipeline before the AWS Cloud Development Kit (AWS CDK) 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 CI/CD pipeline to run the unit testing framework.
C. Install the AWS Serverless Application Model (AWS SAM) CLI tool. Use the sam local generate-event command to generate sample events for the 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 for 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.js base image to invoke the Lambda functions. Check the response. Document how to run the Docker container for the other developers on the team. Update the ClICD pipeline to run the Docker container.

Correct Answer: C
Section:
Explanation:
The AWS Serverless Application Model Command Line Interface (AWS SAM CLI) is a command-line tool for local development and testing of Serverless applications3.Thesam local generate-eventcommand of AWS SAM CLI generates sample events for automated tests3.Thesam local invokecommand is used to invoke Lambda functions3. Therefore, option C is correct.

QUESTION 54
A developer wants 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 test, 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 existing API. 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.

Correct Answer: B
Section:
Explanation:
Amazon API Gateway allows you to create, deploy, and manage a RESTful API to expose backend HTTP endpoints, AWS Lambda functions, or other AWS services1.You can use API Gateway to perform basic validation of an API request before proceeding with the integration request1.When the validation fails, API Gateway immediately fails the request, returns a 400 error response to the caller, and publishes the validation results in CloudWatch Logs1.
To test changes before deploying to a production environment, you can modify the existing API to add request validation and deploy the updated API to a new API Gateway stage1. This allows you to perform tests without affecting the production environment.Once testing is complete and successful, you can then deploy the updated API to the API Gateway production stage1.
This approach has the least operational overhead as it avoids unnecessary creation of new APIs or exporting and importing of APIs.It leverages the existing infrastructure and only requires changes in the configuration of the existing API1.

QUESTION 55
A developer creates a static website for their department The developer deploys the static assets for the website to an Amazon S3 bucket and serves the assets with Amazon CloudFront The developer uses origin access control (OAC) on the CloudFront distribution to access the S3 bucket
The developer notices users can access the root URL and specific pages but cannot access directories without specifying a file name. For example, /products/index.html works, but /products returns an error The developer needs to enable accessing directories without specifying a file name without exposing the S3 bucket publicly.
Which solution will meet these requirements'?

A. Update the CloudFront distribution's settings to index.html as the default root object is set
B. Update the Amazon S3 bucket settings and enable static website hosting. Specify index html as the Index document Update the S3 bucket policy to enable access. Update the CloudFront distribution's origin to use the S3 website endpoint
C. Create a CloudFront function that examines the request URL and appends index.html when directories are being accessed Add the function as a viewer request CloudFront function to the CloudFront distribution's behavior.
D. Create a custom error response on the CloudFront distribution with the HTTP error code set to the HTTP 404 Not Found response code and the response page path to /index html Set the HTTP response code to the HTTP 200 OK response code

Correct Answer: B
Section:
Explanation:
Problem: Directory access without file names fails.
S3 Static Website Hosting:
Configuring S3 as a static website enables automatic serving ofindex.htmlfor directory requests.
Bucket policies ensure correct access permissions.
Updating the CloudFront origin simplifies routing.
Avoiding Public Exposure: The S3 website endpoint allows CloudFront to access content without making the bucket public.
S3 Static Website Hosting:https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html

QUESTION 56
A company needs to deploy all its cloud resources by using AWS CloudFormation templates A developer must create an Amazon Simple Notification Service (Amazon SNS) automatic notification to help enforce this rule. The

developer creates an SNS topic and subscribes the email address of the company's security team to the SNS topic.
The security team must receive a notification immediately if an 1AM role is created without the use of CloudFormation.
Which solution will meet this requirement?

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

Correct Answer: D
Section:
Explanation:
EventBridge (formerly CloudWatch Events) is the ideal service for real-time event monitoring.
CloudTrail logs IAM role creation.
EventBridge rules can filter CloudTrail events and trigger SNS notifications instantly.

QUESTION 57
A developer is working on a web application that uses Amazon DynamoDB as its data store The application has two DynamoDB tables one table that is named artists and one table that is named songs The artists table has artistName as the partition key. The songs table has songName as the partition key and artistName as the sort key
The table usage patterns include the retrieval of multiple songs and artists in a single database operation from the webpage. The developer needs a way to retrieve this information with minimal network traffic and optimal application performance.
Which solution will meet these requirements'?

A. Perform a BatchGetItem operation that returns items from the two tables. Use the list of songName artistName keys for the songs table and the list of artistName key for the artists table.
B. Create a local secondary index (LSI) on the songs table that uses artistName as the partition key Perform a query operation for each artistName on the songs table that filters by the list of songName Perform a query operation for each artistName on the artists table
C. Perform a BatchGetItem operation on the songs table that uses the songName/artistName keys. Perform a BatchGetItem operation on the artists table that uses artistName as the key.
D. Perform a Scan operation on each table that filters by the list of songName/artistName for the songs table and the list of artistName in the artists table.

Correct Answer: A
Section:
Explanation:
Scenario:Application needs to fetch songs and artists efficiently in a single operation.
BatchGetItem:This DynamoDB operation retrieves multiple items across different tables based on their primary keys in a single request.
Optimized for Request Batching:This approach reduces network traffic compared to performing multiple queries individually.
Data Modeling:Thesongstable is designed appropriately for this access pattern usingartistNameas the sort key.
Amazon DynamoDB BatchGetItem:https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.ht

QUESTION 58
A data visualization company wants to strengthen the security of its core applications The applications are deployed on AWS across its development staging, pre-production, and production environments. The company needs to encrypt all of its stored sensitive credentials The sensitive credentials need to be automatically rotated Aversion of the sensitive credentials need to be stored for each environment
Which solution will meet these requirements in the MOST operationally efficient way?

A. Configure AWS Secrets Manager versions to store different copies of the same credentials across multiple environments
B. Create a new parameter version in AWS Systems Manager Parameter Store for each environment Store the environment-specific credentials in the parameter version.

C. Configure the environment variables in the application code Use different names for each environment type
D. Configure AWS Secrets Manager to create a new secret for each environment type. Store the environment-specific credentials in the secret

Correct Answer: D
Section:
Explanation:
Secrets Management:AWS Secrets Manager is designed specifically for storing and managing sensitive credentials.
Environment Isolation:Creating separate secrets for each environment (development, staging, etc.) ensures clear separation and prevents accidental leaks.
Automatic Rotation:Secrets Manager provides built-in rotation capabilities, enhancing security posture.
AWS Secrets Manager:https://aws.amazon.com/secrets-manager/
Secrets Manager Rotation:https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html

QUESTION 59
A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values
When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.
Which solution will meet these requirements with the LEAST development effort?

A. Modify the CloudFormation stack to set the deletion policy to Retain for the Parameter Store parameters.
B. Create an Amazon DynamoDB table as a resource in the CloudFormation stack to hold configuration data for the application Migrate the parameters that the application is modifying from Parameter Store to the DynamoDB table
C. Create an Amazon RDS DB instance as a resource in the CloudFormation stack. Create a table in the database for parameter configuration. Migrate the parameters that the application is modifying from Parameter Store to the configuration table
D. Modify the CloudFormation stack policy to deny updates on Parameter Store parameters

Correct Answer: A
Section:
Explanation:
Problem: CloudFormation updates reset Parameter Store parameters, disrupting application behavior.
Deletion Policy: CloudFormation has a deletion policy that controls resource behavior when a stack is deleted or updated. The 'Retain' policy instructs CloudFormation to preserve a resource's current state.
Least Development Effort: This solution involves a simple CloudFormation template modification, requiring minimal code changes.
CloudFormation Deletion Policies:https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html

QUESTION 60
A company has built an AWS Lambda function to convert large image files into output files that can be used in a third-party viewer application The company recently added a new module to the function to improve the output of the generated files However, the new module has increased the bundle size and has increased the time that is needed to deploy changes to the function code.
How can a developer increase the speed of the Lambda function deployment?

A. Use AWS CodeDeploy to deploy the function code
B. Use Lambda layers to package and load dependencies.
C. Increase the memory size of the function.
D. Use Amazon S3 to host the function dependencies

Correct Answer: B
Section:
Explanation:
Problem: Large bundle size increases Lambda deployment time.

Lambda Layers: Layers let you package dependencies separately from your function code. This optimizes the deployment package, making updates faster.
Modularization: Breaking down dependencies into layers improves code organization and reusability.
AWS Lambda Layers:https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html

QUESTION 61
A company runs a batch processing application by using AWS Lambda functions and Amazon API Gateway APIs with deployment stages for development, user acceptance testing and production A development team needs to configure the APIs in the deployment stages to connect to third-party service endpoints.
Which solution will meet this requirement?

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

Correct Answer: B
Section:
Explanation:
API Gateway Stage Variables: These are designed for configuring dynamic values for your APIs in different deployment stages (dev, test, prod). Here's how to use them for third-party endpoints:
In the API Gateway console,access the 'Stages' section of your API.
For each stage,create a stage variable named something likethirdPartyEndpoint.
Set the value of this variable to the actual endpoint URL for that specific environment.
When configuring API requests within your API Gateway method,reference this endpoint using${stageVariables.thirdPartyEndpoint}.
Why Stage Variables Excel Here:
Environment Isolation:This approach keeps the endpoint configuration specific to each deployment stage,ensuring the right endpoints are used during development,testing,and production cycles.
Ease of Management:You manage the endpoints directly through the API Gateway console without additional infrastructure.
Amazon API Gateway Stage Variables:https://docs.aws.amazon.com/apigateway/latest/developerguide/stage-variables.html

QUESTION 62
A developer is investigating an issue in part of a company's application. In the application messages are sent to an Amazon Simple Queue Service (Amazon SQS) queue The AWS Lambda function polls messages from the SQS queue and sends email messages by using Amazon Simple Email Service (Amazon SES) Users have been receiving duplicate email messages during periods of high traffic.
Which reasons could explain the duplicate email messages? (Select TWO.)

A. Standard SQS queues support at-least-once message delivery
B. Standard SQS queues support exactly-once processing, so the duplicate email messages are because of user error.
C. Amazon SES has the DomainKeys Identified Mail (DKIM) authentication incorrectly configured
D. The SQS queue's visibility timeout is lower than or the same as the Lambda function's timeout.
E. The Amazon SES bounce rate metric is too high.

Correct Answer: A
Section:
Explanation:
SQS Delivery Behavior:Standard SQS queues guarantee at-least-once delivery, meaning messages may be processed more than once. This can lead to duplicate emails in this scenario.
Visibility Timeout:If the visibility timeout on the SQS queue is too short, a message might become visible for another consumer before the first Lambda function finishes processing it. This can also lead to duplicates.
Amazon SQS Delivery Semantics:[invalid URL removed]
Amazon SQS Visibility Timeout:https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html

QUESTION 63
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 64
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 65
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 66
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 67
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 Deleteltem 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 Deleteltem 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 68
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 69
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 70
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 71
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 72
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 73
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 74
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 75
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 76
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 77
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 78
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 79
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 80
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 81
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 82
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 83
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 84
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 85
A developer is building a microservices-based application by using Python on AWS and several AWS services The developer must use AWS X-Ray The developer views the service map by using the console to view the service dependencies. During testing, the developer notices that some services are missing from the service map
What can the developer do to ensure that all services appear in the X-Ray service map?

A. Modify the X-Ray Python agent configuration in each service to increase the sampling rate
B. Instrument the application by using the X-Ray SDK for Python. Install the X-Ray SDK for all the services that the application uses
C. Enable X-Ray data aggregation in Amazon CloudWatch Logs for all the services that the application uses
D. Increase the X-Ray service map timeout value in the X-Ray console

Correct Answer: B
Section:
Explanation:

AWS X-Ray SDK:The primary way to enable X-Ray tracing within applications. The SDK sends data about requests and subsegments to the X-Ray daemon for service map generation.
Instrumenting All Services:To visualize a complete microservice architecture on the service map, each relevant service must include the X-Ray SDK.
AWS X-Ray Documentation:https://docs.aws.amazon.com/xray/
X-Ray SDK for Python:https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-python.html

QUESTION 86
A developer is creating an AWS Lambda function. The Lambda function needs an external library to connect to a third-party solution The external library is a collection of files with a total size of 100 MB The developer needs to make the external library available to the Lambda execution environment and reduce the Lambda package space
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a Lambda layer to store the external library Configure the Lambda function to use the layer
B. Create an Amazon S3 bucket Upload the external library into the S3 bucket. Mount the S3 bucket folder in the Lambda function Import the library by using the proper folder in the mount point.
C. Load the external library to the Lambda function's /tmp directory during deployment of the Lambda package. Import the library from the /tmp directory.
D. Create an Amazon Elastic File System (Amazon EFS) volume. Upload the external library to the EFS volume Mount the EFS volume in the Lambda function. Import the library by using the proper folder in the mount point.

Correct Answer: A
Section:
Explanation:
Lambda Layers:These are designed to package dependencies that you can share across functions.
How to Use:
Create a layer, upload your 100MB library as a zip.
Attach the layer to your function.
In your function code, import the library from the standard layer path.
Lambda Layers:https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html

QUESTION 87
A company built an online event platform For each event the company organizes quizzes and generates leaderboards that are based on the quiz scores. The company stores the leaderboard data in Amazon DynamoDB and retains the data for 30 days after an event is complete The company then uses a scheduled job to delete the old leaderboard data
The DynamoDB table is configured with a fixed write capacity. During the months when many events occur, the DynamoDB write API requests are throttled when the scheduled delete job runs.
A developer must create a long-term solution that deletes the old leaderboard data and optimizes write throughput
Which solution meets these requirements?

A. Configure a TTL attribute for the leaderboard data
B. Use DynamoDB Streams to schedule and delete the leaderboard data
C. Use AWS Step Functions to schedule and delete the leaderboard data.
D. Set a higher write capacity when the scheduled delete job runs

Correct Answer: A
Section:
Explanation:
DynamoDB TTL (Time-to-Live):A native feature that automatically deletes items after a specified expiration time.
Efficiency:Eliminates the need for scheduled deletion jobs, optimizing write throughput by avoiding potential throttling conflicts.
Seamless Integration:TTL works directly within DynamoDB, requiring minimal development overhead.
DynamoDB TTL Documentation:https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html

QUESTION 88
A developer must use multi-factor authentication (MFA) to access data in an Amazon S3 bucket that is in another AWS account. Which AWS Security Token Service (AWS STS) API operation should the developer use with the MFA information to meet this requirement?

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

Correct Answer: D
Section:
Explanation:
AWS STS AssumeRole:The central operation for assuming temporary security credentials, commonly used for cross-account access.
MFA Integration:TheAssumeRolecall can include MFA information to enforce multi-factor authentication.
Credentials for S3 Access:The returned temporary credentials would provide the necessary permissions to access the S3 bucket in the other account.
AWS STS AssumeRole Documentation:https://docs.aws.amazon.com/STS/latest/APIReference/API_AssumeRole.html

QUESTION 89
A company has an analytics application that uses an AWS Lambda function to process transaction data asynchronously A developer notices that asynchronous invocations of the Lambda function sometimes fail When failed Lambda function invocations occur, the developer wants to invoke a second Lambda function to handle errors and log details.
Which solution will meet these requirements?

A. Configure a Lambda function destination with a failure condition Specify Lambda function as the destination type Specify the error-handling Lambda function's Amazon Resource Name (ARN) as the resource
B. Enable AWS X-Ray active tracing on the initial Lambda function. Configure X-Ray to capture stack traces of the failed invocations. Invoke the error-handling Lambda function by including the stack traces in the event object.
C. Configure a Lambda function trigger with a failure condition Specify Lambda function as the destination type Specify the error-handling Lambda function's Amazon Resource Name (ARN) as the resource
D. Create a status check alarm on the initial Lambda function. Configure the alarm to invoke the error-handling Lambda function when the alarm is initiated. Ensure that the alarm passes the stack trace in the event object.

Correct Answer: A
Section:
Explanation:
Lambda Destinations on Failure:Allow routing asynchronous function invocations to specified resources (like another Lambda function) upon failure.
Error Handling:The error-handling Lambda receives details about the failure, enabling logging and custom actions.
Direct Integration:This solution leverages native Lambda functionality for a simpler implementation.

QUESTION 90
A company is preparing to migrate an application to the company's first AWS environment Before this migration, a developer is creating a proof-of-concept application to validate a model for building and deploying container-based applications on AWS.
Which combination of steps should the developer take to deploy the containerized proof-of-concept application with the LEAST operational effort? (Select TWO.)

A. Package the application into a zip file by using a command line tool Upload the package to Amazon S3
B. Package the application into a container image by using the Docker CLI. Upload the image to Amazon Elastic Container Registry (Amazon ECR)
C. Deploy the application to an Amazon EC2 instance by using AWS CodeDeploy.
D. Deploy the application to Amazon Elastic Kubernetes Service (Amazon EKS) on AWS Fargate
E. Deploy the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate

Correct Answer: B, E
Section:
Explanation:
Containerization:Packaging the application as a container image promotes portability and standardization. Docker is the standard tool for containerization.
Amazon ECR:ECR is a managed container registry designed to work seamlessly with AWS container services.
Fargate:ECS Fargate provides serverless container orchestration, minimizing operational overhead for this proof-of-concept.

Docker:https://www.docker.com/
Amazon ECR:https://aws.amazon.com/ecr/

QUESTION 91
A company runs an application on AWS The application stores data in an Amazon DynamoDB table Some queries are taking a long time to run These slow queries involve an attribute that is not the table's partition key or sort key
The amount of data that the application stores in the DynamoDB table is expected to increase significantly. A developer must increase the performance of the queries.
Which solution will meet these requirements'?

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

Correct Answer: B
Section:
Explanation:
Global Secondary Index (GSI):GSIs enable alternative query patterns on a DynamoDB table by using different partition and sort keys.
Addressing Query Bottleneck:By making the slow-query attribute the GSI's partition key, you optimize queries on that attribute.
Scalability:GSIs automatically scale to handle increasing data volumes.
Amazon DynamoDB Global Secondary Indexes:https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html

QUESTION 92
A developer maintains a critical business application that uses Amazon DynamoDB as the primary data store The DynamoDB table contains millions of documents and receives 30-60 requests each minute The developer needs to perform processing in near-real time on the documents when they are added or updated in the DynamoDB table
How can the developer implement this feature with the LEAST amount of change to the existing application code?

A. Set up a cron job on an Amazon EC2 instance Run a script every hour to query the table for changes and process the documents
B. Enable a DynamoDB stream on the table Invoke an AWS Lambda function to process the documents.
C. Update the application to send a PutEvents request to Amazon EventBridge. Create an EventBridge rule to invoke an AWS Lambda function to process the documents.
D. Update the application to synchronously process the documents directly after the DynamoDB write

Correct Answer: B
Section:
Explanation:
DynamoDB Streams:Capture near real-time changes to DynamoDB tables, triggering downstream actions.
Lambda for Processing:Lambda functions provide a serverless way to execute code in response to events like DynamoDB Stream updates.
Minimal Code Changes:This solution requires the least modifications to the existing application.
DynamoDB Streams:https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html
AWS Lambda:https://aws.amazon.com/lambda/

QUESTION 93
A team is developing an application that is deployed on Amazon EC2 instances. During testing, the team receives an error. The EC2 instances are unable to access an Amazon S3 bucket.
Which steps should the team take to troubleshoot this issue? (Select TWO.)

A. Check whether the policy that is assigned to the JAM role that is attached to the EC2 instances grants access to Amazon S3.
B. Check the S3 bucket policy to validate the access permissions for the S3 bucket.
C. Check whether the policy that is assigned to the 1AM user that is attached to the EC2 instances grants access to Amazon S3.

D. Check the S3 Lifecycle policy to validate the permissions that are assigned to the S3 bucket.
E. Check the security groups that are assigned to the EC2 instances. Make sure that a rule is not blocking the access to Amazon S3.

Correct Answer: A, B
Section:

QUESTION 94
A developer created several AWS Lambda functions that write data to a single Amazon S3 bucket. The developer configured all the Lambda functions to send logs and metrics to Amazon CloudWatch.
The developer receives reports that one of the Lambda functions writes data to the bucket very slowly. The developer needs to measure the latency between the problematic Lambda function and the S3 bucket.
Which solution will meet this requirement?

A. Enable AWS X-Ray on the Lambda function. In the generated trace map. select the line between Lambda and Amazon S3.
B. Query the Lambda function's log file in Amazon CloudWatch Logs Insights. Return the average of the auto-discovered duration field.
C. Enable CloudWatch Lambda Insights on the function. View the latency graph that CloudWatch Lambda Insights provides.
D. Enable AWS X-Ray on the Lambda function. Select Amazon S3 in the latency graph to view the latency histogram.

Correct Answer: A
Section:

QUESTION 95
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 96
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 97
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 98
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 99
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: