# AWS-Certified-Developer-Associate Dumps

# Amazon AWS Certified Developer - Associate

# https://www.certleader.com/AWS-Certified-Developer-Associate-dumps.html

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

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

**Answer:** D

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

**NEW QUESTION 2**
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 queu
B. Configure the visibility timeout value.
C. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ
D. Configure the DelaySeconds values.
E. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ
F. Configure the visibility timeout value.
G. Use an SQS FIFO queu
H. Configure the DelaySeconds value.

**Answer:** A

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

**NEW QUESTION 3**
A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node is application. To minimize these bugs, the developer wants to impendent automated testing of Lambda functions in an environment that Closely simulates the Lambda environment.
The developer needs to give other developers the ability to run the tests locally. The developer also needs to integrate the tests into the team's continuous integration and continuous delivery (Ct/CO) pipeline before the AWS Cloud Development Kit (AWS COK) deployment.
Which solution will meet these requirements?

A. Create sample events based on the Lambda documentatio
B. Create automated test scripts that use the cdk local invoke command to invoke the Lambda function
C. 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. Install a unit testing framework that reproduces the Lambda execution environmen
E. Create sample events based on the Lambda Documentation Invoke the handler function by using a unit testing framewor
F. Check the response Document how to run the unit testing framework for the other developers on the tea
G. Update the OCD pipeline to run the unit testing framework.
H. Install the AWS Serverless Application Model (AWS SAW) CLI tool Use the Sam local generate-event command to generate sample events for me automated test
I. Create automated test scripts that use the Sam local invoke command to invoke the Lambda function
J. Check the response Document the test scripts tor the other developers on the team Update the CI/CD pipeline to run the test scripts.
K. Create sample events based on the Lambda documentatio

L. Create a Docker container from the Node is base image to invoke the Lambda function
M. Check the response Document how to run the Docker container for the more developers on the team update the CI/CD pipeline to run the Docker container.

**Answer:** C

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


**NEW QUESTION 4**
A company uses Amazon API Gateway to expose a set of APIs to customers. The APIs have caching enabled in API Gateway. Customers need a way to invalidate the cache for each API when they test the API.
What should a developer do to give customers the ability to invalidate the API cache?

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

**Answer:** D


**NEW QUESTION 5**
                            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.

**Answer:** B

**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


**NEW QUESTION 6**
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 data. 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 instanc
B. Enable encryption of data in transit and at res
C. Store frequently accessed data in the cache.
D. Create an Amazon ElastiCache for Memcached instanc
E. Enable encryption of data in transit and at res
F. Store frequently accessed data in the cache.
G. Create an Amazon RDS for MySQL read replic
H. Connect to the read replica by using SS
I. Configure the read replica to store frequently accessed data.
J. Create an Amazon DynamoDB table and a DynamoDB Accelerator (DAX) cluster for the tabl
K. Store frequently accessed data in the DynamoDB table.

**Answer:** A

**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.
References:
? [What Is Amazon ElastiCache? - Amazon ElastiCache]
                            ? [Encryption in Transit - Amazon ElastiCache for Redis]
? [Encryption at Rest - Amazon ElastiCache for Redis]

**NEW QUESTION 7**
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. Mastered
B. Not Mastered

**Answer:** A

**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.
References:
? [Amazon Machine Images (AMI) - Amazon Elastic Compute Cloud]
? [Encrypting an Amazon EBS Snapshot - Amazon Elastic Compute Cloud]
? [Copying an AMI - Amazon Elastic Compute Cloud]

**NEW QUESTION 8**
A company needs to harden its container images before the images are in a running state. The company's application uses Amazon Elastic Container Registry (Amazon ECR) as an image registry. Amazon Elastic Kubernetes Service (Amazon EKS) for compute, and an AWS CodePipeline pipeline that orchestrates a continuous integration and continuous delivery (CI/CD) workflow.
Dynamic application security testing occurs in the final stage of the pipeline after a new image is deployed to a development namespace in the EKS cluster. A developer needs to
                          place an analysis stage before this deployment to analyze the container image earlier in the CI/CD pipeline.
Which solution will meet these requirements with the MOST operational efficiency?

A. Build the container image and run the docker scan command locall
B. Mitigate any findings before pushing changes to the source code repositor
C. Write a pre-commit hook that enforces the use of this workflow before commit.
D. Create a new CodePipeline stage that occurs after the container image is buil
E. Configure ECR basic image scanning to scan on image pus
F. Use an AWS Lambda function as the action provide
G. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings.
H. Create a new CodePipeline stage that occurs after source code has been retrieved from its repositor
I. Run a security scanner on the latest revision of the source cod
J. Fail the pipeline if there are findings.
K. Add an action to the deployment stage of the pipeline so that the action occurs before the deployment to the EKS cluste
L. Configure ECR basic image scanning to scan on image pus
M. Use an AWS Lambda function as the action provide
N. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings.

**Answer:** B

**Explanation:**
The solution that will meet the requirements with the most operational efficiency is to create a new CodePipeline stage that occurs after the container image is built. Configure ECR basic image scanning to scan on image push. Use an AWS Lambda function as the action provider. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings. This way, the container image is analyzed earlier in the CI/CD pipeline and any vulnerabilities are detected and reported before deploying to the EKS cluster. The other options either delay the analysis until after deployment, which increases the risk of exposing insecure images, or perform analysis on the source code instead of the container image, which may not capture all the dependencies and configurations that affect the security posture of the image.
Reference: Amazon ECR image scanning

**NEW QUESTION 9**
A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream. The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is increasing and that the Lambda run duration is constantly above normal.
Which actions should the developer take to increase the processing speed? (Choose two.)

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

**Answer:** AC

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

**NEW QUESTION 10**
A company uses a custom root certificate authority certificate chain (Root CA Cert) that is 10 KB in size generate SSL certificates for its on-premises HTTPS

endpoints. One of the company's cloud based applications has hundreds of AWS Lambda functions that pull date from these endpoints. A developer updated the trust store of the Lambda execution environment to use the Root CA Cert when the Lambda execution environment is initialized. The developer bundled the Root CA Cert as a text file in the Lambdas deployment bundle.

After 3 months of development the root CA Cert is no longer valid and must be updated. The developer needs a more efficient solution to update the Root CA Cert for all deployed Lambda functions. The solution must not include rebuilding or updating all Lambda functions that use the Root CA Cert. The solution must also work for all development, testing and production environment. Each environment is managed in a separate AWS account.

When combination of steps Would the developer take to meet these environments MOST cost-effectively? (Select TWO)

A. Mastered
B. Not Mastered

**Answer:** A

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

**NEW QUESTION 10**
A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetObject",
                "s3:PutObject"
            ],
            "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
        },
        {
            "Effect": "Deny",
            "Action": "s3:*",
            "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/secrets*"
        }
    ]
}
```

Which access does the policy allow regarding the s3:GetObject and s3:PutObject actions?

A. Access on all buckets except the "DOC-EXAMPLE-BUCKET" bucket

B. Access on all buckets that start with "DOC-EXAMPLE-BUCKET" except the "DOC-EXAMPLE-BUCKET/secrets" bucket
C. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket along with access to all S3 actions for objects in the "DOC-EXAMPLE-BUCKET" bucket that start with "secrets"
D. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket except on objects that start with "secrets"

**Answer:** D

**Explanation:**
 The IAM policy shown in the image is a resource-based policy that grants or denies access to an S3 bucket based on certain conditions. The first statement allows access to any S3 action on any object in the "DOC-EXAMPLE-BUCKET" bucket when the request is made over HTTPS (the value of aws:SecureTransport is true). The second statement denies access to the s3:GetObject and s3:PutObject actions on any object in the "DOC-EXAMPLE-BUCKET/secrets" prefix when the request is made over HTTP (the value of aws:SecureTransport is false). Therefore, the policy allows access on all objects in the "DOC-EXAMPLE-BUCKET" bucket except on objects that start with "secrets".
Reference: Using IAM policies for Amazon S3

**NEW QUESTION 12**
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 instanc
B. Store the unique identifier for each request in a database tabl
C. Modify the Lambda function to check the table for the identifier before processing the request.

D. Create an Amazon DynamoDB tabl
E. Store the unique identifier for each request in the tabl
F. Modify the Lambda function to check the table for the identifier before processing the request.
G. Create an Amazon DynamoDB tabl
H. Store the unique identifier for each request in the tabl
I. Modify the Lambda function to return a client error response when the function receives a duplicate request.
J. Create an Amazon ElastiCache for Memcached instanc
K. Store the unique identifier for each request in the cach
L. Modify the Lambda function to check the cache for the identifier before processing the request.

**Answer:** B

**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

**NEW QUESTION 13**
A company has an application that runs as a series of AWS Lambda functions. Each Lambda function receives data from an Amazon Simple Notification Service (Amazon SNS) topic and writes the data to an Amazon Aurora DB instance.
To comply with an information security policy, the company must ensure that the Lambda functions all use a single securely encrypted database connection string to access Aurora.
Which solution will meet these requirements'?

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

**Answer:** A

**Explanation:**

 This solution will meet the requirements by using IAM database authentication for Aurora which enables using IAM roles or users to authenticate with Aurora databases instead of using passwords or other secrets. The developer can use IAM database authentication for Aurora to enable secure database connections for all the Lambda functions that access Aurora DB instance. The developer can create an IAM role with permission to connect to Aurora DB instance and attach it to each Lambda function. The developer can also configure Aurora DB instance to use IAM database authentication and enable encryption in transit using SSL certificates. This way, the Lambda functions can use a single securely encrypted database connection string to access Aurora without needing any secrets or passwords. Option B is not optimal because it will store the credentials and read them from an encrypted Amazon RDS DB instance, which may introduce additional costs and complexity for managing and accessing another RDS DB instance. Option C is not optimal because it will store the credentials in AWS Systems Manager Parameter Store as a secure string parameter, which may require additional steps or permissions to retrieve and decrypt the credentials from Parameter Store. Option D is not optimal because it will use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption, which may not be secure or scalable as environment variables are stored as plain text unless encrypted with AWS KMS. References: [IAM Database Authentication for MySQL and PostgreSQL], [Using SSL/TLS to Encrypt a Connection to a DB Instance]

**NEW QUESTION 17**
A developer must use multi-factor authentication (MFA) to access data in an Amazon S3 bucket that is in another AWS account. Which AWS Security Token Service (AWS STS) API operation should the developer use with the MFA information to meet this requirement?

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 18**
For a deployment using AWS Code Deploy, what is the run order of the hooks for in-place deployments?

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

**Answer:** B

**Explanation:**
For in-place deployments, AWS CodeDeploy uses a set of predefined hooks that run in a specific order during each deployment lifecycle event. The hooks are ApplicationStop, BeforeInstall, AfterInstall, ApplicationStart, and ValidateService. The run order of the hooks for in-place deployments is as follows:
? ApplicationStop: This hook runs first on all instances and stops the current application that is running on the instances.

? BeforeInstall: This hook runs after ApplicationStop on all instances and performs any tasks required before installing the new application revision.
? AfterInstall: This hook runs after BeforeInstall on all instances and performs any
tasks required after installing the new application revision.
? ApplicationStart: This hook runs after AfterInstall on all instances and starts the new application that has been installed on the instances.
? ValidateService: This hook runs last on all instances and verifies that the new application is running properly on the instances.
Reference: [AWS CodeDeploy lifecycle event hooks reference]


**NEW QUESTION 22**
An developer is building a serverless application by using the AWS Serverless Application Model (AWS SAM). The developer is currently testing the application in a development environment. When the application is nearly finsihed, the developer will need to set up
additional testing and staging environments for a quality assurance team.
The developer wants to use a feature of the AWS SAM to set up deployments to multiple environments.
Which solution will meet these requirements with the LEAST development effort?

A. Add a configuration file in TOML format to group configuration entries to every environmen
B. Add a table for each testing and staging environmen
C. Deploy updates to the environments by using the sam deploy command and the --config-env flag that corresponds to the each environment.
D. Create additional AWS SAM templates for each testing and staging environmen
E. Write a custom shell script that uses the sam deploy command and the --template-file flag to deploy updates to the environments.
F. Create one AWS SAM configuration file that has default parameter
G. Perform updates to the testing and staging environments by using the —parameter-overrides flag in the AWS SAM CLI and the parameters that the updates will override.
H. Use the existing AWS SAM templat
I. Add additional parameters to configure specific attributes for the serverless function and database table resources that are in each environmen
J. Deploy updates to the testing and staging environments by using the sam deploy command.

**Answer:** A

**Explanation:**
The correct answer is A. Add a configuration file in TOML format to group configuration entries to every environment. Add a table for each testing and staging environment. Deploy updates to the environments by using the sam deploy command and the --config-env flag that corresponds to the each environment.
* A. Add a configuration file in TOML format to group configuration entries to every environment. Add a table for each testing and staging environment. Deploy updates to the environments by using the sam deploy command and the --config-env flag that corresponds to the each environment. This is correct. This solution will meet the requirements with the least development effort, because it uses a feature of the AWS SAM CLI that supports a project-level configuration file that can be used to configure AWS SAM CLI command parameter values1. The configuration file can have multiple environments, each with its own set of parameter values, such as stack name, region, capabilities, and more2. The developer can use the --config-env option to specify which environment to use when deploying the application3. This way, the developer can avoid creating multiple templates or scripts, or manually overriding parameters for each environment.
* B. Create additional AWS SAM templates for each testing and staging environment. Write a custom shell script that uses the sam deploy command and the --template-file flag to
deploy updates to the environments. This is incorrect. This solution will not meet the requirements with the least development effort, because it requires creating and maintaining multiple templates and scripts for each environment. This can introduce duplication, inconsistency, and complexity in the deployment process.
* C. Create one AWS SAM configuration file that has default parameters. Perform updates to the testing and staging environments by using the —parameter-overrides flag in the AWS SAM CLI and the parameters that the updates will override. This is incorrect. This solution will not meet the requirements with the least development effort, because it requires manually specifying and overriding parameters for each environment every time the developer deploys the application. This can be error-prone, tedious, and inefficient.
* D. Use the existing AWS SAM template. Add additional parameters to configure specific attributes for the serverless function and database table resources that are in each environment. Deploy updates to the testing and staging environments by using the sam deploy command. This is incorrect. This solution will not meet the requirements with the least development effort, because it requires modifying the existing template and adding complexity to the resource definitions for each environment. This can also make it difficult to manage and track changes across different environments.
References:
? 1: AWS SAM CLI configuration file - AWS Serverless Application Model
? 2: Configuration file basics - AWS Serverless Application Model
? 3: Specify a configuration file - AWS Serverless Application Model


**NEW QUESTION 24**
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 documentatio
B. Create automated test scripts that use the cdk local invoke command to invoke the Lambda function
C. Check the respons
D. Document the test scripts for the other developers on the tea
E. Update the CI/CD pipeline to run the test scripts.
F. Install a unit testing framework that reproduces the Lambda execution environment.
G. Invoke the handler function by using a unit testing framewor
H. Check the respons
I. Document how to run the unit testing framework for the other developers on the tea
J. Update the CI/CD pipeline to run the unit testing framework.
K. Install the AWS Serverless Application Model (AWS SAM) CLI too
L. Use the sam local generate-event command to generate sample events for the automated test
M. Create automated test scripts that use the sam local invoke command to invoke the Lambda function
N. Check the respons
O. Document the test scripts for the other developers on the tea
P. Update the CI/CD pipeline to run the test scripts.
Q. Create sample events based on the Lambda documentatio

R. Create a Docker container from the Node.js base image to invoke the Lambda function
S. Check the respons
T. Document how to run the Docker container for the other developers on the tea
. Update the CIICD pipeline to run the Docker container.

**Answer:** C

**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. The sam local generate-event command of AWS SAM CLI generates sample events for automated tests3. The sam local invoke command is used to invoke Lambda functions3. Therefore, option C is correct.

---

**NEW QUESTION 28**
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions. When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

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

**Answer:** A

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

---

**NEW QUESTION 33**
A developer at a company needs to create a small application that makes the same API call once each day 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 Kubernetes cron job that runs on Amazon Elastic Kubernetes Service (Amazon EKS).
B. Use an Amazon Linux crontab scheduled job that runs on Amazon EC2.
C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event.
D. Use an AWS Batch job that is submitted to an AWS Batch job queue.

**Answer:** C

**Explanation:**
 The correct answer is C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event.
* C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event. This is correct. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda runs your code on a high-availability compute infrastructure and performs all of the administration of the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, and logging1. Amazon EventBridge is a serverless event bus service that enables you to connect your applications with data from a variety of sources2. EventBridge can create rules that run on a schedule, either at regular intervals or at specific times and dates, and invoke targets such as Lambda functions3. This solution meets the requirements of creating a small application that makes the same API call once each day at a designated time, without requiring any infrastructure in the AWS Cloud or any operational overhead.
* A. Use a Kubernetes cron job that runs on Amazon Elastic Kubernetes Service (Amazon EKS). This is incorrect. Amazon EKS is a fully managed Kubernetes service that allows you to run containerized applications on AWS4. Kubernetes cron jobs are tasks that run periodically on a given schedule5. This solution could meet the functional requirements of creating a small application that makes the same API call once each day at a designated time, but it would not be the most operationally efficient manner. The company would need to provision and manage an EKS cluster, which would incur additional costs and complexity.
* B. Use an Amazon Linux crontab scheduled job that runs on Amazon EC2. This is incorrect. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud6. Crontab is a Linux utility that allows you to schedule commands or scripts to run automatically at a specified time or date7. This solution could meet the functional requirements of creating a small application that makes the same API call once each day at a designated time, but it would not be the most operationally efficient manner. The company would need to provision and manage an EC2 instance, which would incur additional costs and complexity.
* D. Use an AWS Batch job that is submitted to an AWS Batch job queue. This is incorrect. AWS Batch enables you to run batch computing workloads on the AWS Cloud8. Batch jobs are units of work that can be submitted to job queues, where they are executed in parallel or sequentially on compute environments9. This solution could meet the functional requirements of creating a small application that makes the same API call once each day at a designated time, but it would not be the most operationally efficient manner. The company would need to configure and manage an AWS Batch environment, which would incur additional costs and complexity.
References:
? 1: What is AWS Lambda? - AWS Lambda
? 2: What is Amazon EventBridge? - Amazon EventBridge
? 3: Creating an Amazon EventBridge rule that runs on a schedule - Amazon EventBridge
? 4: What is Amazon EKS? - Amazon EKS
? 5: CronJob - Kubernetes
? 6: What is Amazon EC2? - Amazon EC2
? 7: Crontab in Linux with 20 Useful Examples to Schedule Jobs - Tecmint
? 8: What is AWS Batch? - AWS Batch

? 9: Jobs - AWS Batch

**NEW QUESTION 35**
A developer is building a new application on AWS. The application uses an AWS Lambda function that retrieves information from an Amazon DynamoDB table. The developer hard coded the DynamoDB table name into the Lambda function code. The table name might change over time. The developer does not want to modify the Lambda code if the table name changes.
Which solution will meet these requirements MOST efficiently?

A. Create a Lambda environment variable to store the table nam
B. Use the standard method for the programming language to retrieve the variable.
C. Store the table name in a fil
D. Store the file in the /tmp folde
E. Use the SDK for the programming language to retrieve the table name.
F. Create a file to store the table nam
G. Zip the file and upload the file to the Lambda laye
H. Use the SDK for the programming language to retrieve the table name.
                        Create a global variable that is outside the handler in the Lambda function to store the table name.
I.

**Answer:** A

**Explanation:**
The solution that will meet the requirements most efficiently is to create a Lambda environment variable to store the table name. Use the standard method for the programming language to retrieve the variable. This way, the developer can avoid hard-coding the table name in the Lambda function code and easily change the table name by updating the environment variable. The other options either involve storing the table name in a file, which is less efficient and secure than using an environment variable, or creating a global variable, which is not recommended as it can cause concurrency issues.
Reference: Using AWS Lambda environment variables

**NEW QUESTION 37**
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 use
D. Add the user access key ID and secret access key as environment variables in the Lambda function.
E. Add the AWS account root user access key ID and secret access key as encrypted environment variables in the Lambda function

**Answer:** B

**Explanation:**
The most secure way to grant the Lambda function access to the S3 bucket and the DynamoDB table is to create an IAM role for the Lambda function and attach the existing IAM policy to the role. This way, you can use the principle of least privilege and avoid exposing any credentials in your function code or environment variables. You can also leverage the temporary security credentials that AWS provides to the Lambda function when it assumes the role. This solution follows the best practices for working with AWS Lambda functions1 and designing and architecting with DynamoDB2. References
? Best practices for working with AWS Lambda functions
? Best practices for designing and architecting with DynamoDB

**NEW QUESTION 41**
A developer has written an AWS Lambda function. The function is CPU-bound. The developer wants to ensure that the function returns responses quickly.
How can the developer improve the function's performance?

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

**Answer:** B

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

**NEW QUESTION 45**
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

**Answer:** C

**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

## NEW QUESTION 50

A company hosts a client-side web application for one of its subsidiaries on Amazon S3. The web application can be accessed through Amazon CloudFront from https://www.example.com. After a successful rollout, the company wants to host three more client-side web applications for its remaining subsidiaries on three separate S3 buckets.

To achieve this goal, a developer moves all the common JavaScript files and web fonts to a central S3 bucket that serves the web applications. However, during testing, the developer notices that the browser blocks the JavaScript files and web fonts.

What should the developer do to prevent the browser from blocking the JavaScript files and web fonts?

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

**Answer:** C

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

## NEW QUESTION 52

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

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

**Answer:** D

**Explanation:**
This solution will meet the requirements by creating a single S3 bucket policy that denies access to the S3 bucket unless the request comes from one of the specified VPC endpoints. The aws:SourceVpce condition key is used to match the ID of the VPC endpoint that is used to access the S3 bucket. The StringNotEquals condition operator is used to negate the condition, so that only requests from the listed VPC endpoints are allowed.
Option A is not optimal because it will create multiple S3 bucket policies, which is not possible as only one bucket policy can be attached to an S3 bucket. Option B is not optimal because it will use the aws:SourceVpc condition key, which matches the ID of the VPC that is used to access the S3 bucket, not the VPC endpoint. Option C is not optimal because it will use the StringNotEquals condition operator with a single value, which will deny access to the S3 bucket from all VPC endpoints except one.
References: Using Amazon S3 Bucket Policies and User Policies, AWS Global Condition Context Keys

## NEW QUESTION 55

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.

**Answer:** C

**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. References:
? 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

**NEW QUESTION 59**
A company has installed smart motes in all Its customer locations. The smart meter's measure power usage at 1minute intervals and send the usage readings to a remote endpoint tot collection. The company needs to create an endpoint that will receive the smart meter readings and store the readings in a database. The company wants to store the location ID and timestamp information.
The company wants to give Is customers low-latency access to their current usage and historical usage on demand The company expects demand to increase significantly. The solution must not impact performance or include downtime write seeing.
When solution will meet these requirements MOST cost-effectively?

A. Store the smart meter readings in an Amazon RDS databas
B. Create an index on the location ID and timestamp columns Use the columns to filter on the customers 'data.
C. Store the smart motor readings m an Amazon DynamoDB table Croato a composite Key oy using the location ID and timestamp column
D. Use the columns to filter on the customers' data.
E. Store the smart meter readings in Amazon EastCache for Reds Create a Sorted set key y using the location ID and timestamp column
F. Use the columns to filter on the customers' data.
G. Store the smart meter readings m Amazon S3 Parton the data by using the location ID and timestamp column
H. Use Amazon Athena lo tiler on me customers' data.

**Answer:** B

**Explanation:**
 The solution that will meet the requirements most cost-effectively is to store the smart meter readings in an Amazon DynamoDB table. Create a composite key by using the location ID and timestamp columns. Use the columns to filter on the customers' data. This way, the company can leverage the scalability, performance, and low latency of DynamoDB to store and retrieve the smart meter readings. The company can also use the composite key to query the data by location ID and timestamp efficiently. The other options either involve more expensive or less scalable services, or do not provide low-latency access to the current usage.
Reference: Working with Queries in DynamoDB

**NEW QUESTION 60**
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

A. All at once
B. Rolling with additional batch
C. Bluegreen
D. Immutable

**Answer:** B

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

**NEW QUESTION 65**
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

**Answer:** A

**Explanation:**

"deletes the item from your table without consuming any write throughput" https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html

**NEW QUESTION 66**
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions.
When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

A. Set the Deployment Preference Type to Canary10Percent10Minute
B. Set the                            AutoPublishAlias property to the Lambda alias.
C. Set the Deployment Preference Type to Linear1OPercentEvery10Minute
D. Set AutoPublishAlias property to the Lambda alias.
E. Set the Deployment Preference Type to Canary1OPercent1OMinute
F. Set the PreTraffic and PostTraffic properties to the Lambda alias.
G. Set the Deployment Preference Type to Linear1OPercentEvery1OMinute
H. Set PreTraffic and Post Traffic properties to the Lambda alias.

**Answer:** A

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

**NEW QUESTION 70**
A company is running Amazon EC2 instances in multiple AWS accounts. A developer needs to implement an application that collects all the lifecycle events of the EC2 instances. The application needs to store the lifecycle events in a single Amazon Simple Queue Service (Amazon SQS) queue in the company's main AWS account for further processing.
Which solution will meet these requirements?

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

**Answer:** D

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

**NEW QUESTION 72**
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 backen
B. Conned the repository branches that correspond to each of the desired environment
C. Start deployments by merging code changes to a desired branch.
D. Host each website in AWS Elastic Beanstalk with multiple environment
E. Use the EB CLI to link each repository branc
F. Integrate AWS CodePipeline to automate deployments from version control code merges.
G. Host each website in different Amazon S3 buckets for each environmen
H. Configure AWS CodePipeline to pull source code from version contro
I. Add an AWS CodeBuild stage to copy source code to Amazon S3.
J. Host each website on its own Amazon EC2 instanc
K. Write a custom deployment script to bundle each website's static asset
L. Copy the assets to Amazon EC2. Set up a workflow to run the script when code is merged.

**Answer:** A

**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

**NEW QUESTION 76**
A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values
When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.
Which solution will meet these requirements with the LEAST development effort?

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

**Answer:** D

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

**NEW QUESTION 77**
A team of developed is using an AWS CodePipeline pipeline as a continuous integration and continuous delivery (CI/CD) mechanism for a web application. A developer has written unit tests to programmatically test the functionality of the application code. The unit tests produce a test report that shows the results of each individual check. The developer now
wants to run these tests automatically during the CI/CD process.

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

**Answer:** C

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

**NEW QUESTION 82**
A company is expanding the compatibility of its photo-snaring mobile app to hundreds of additional devices with unique screen dimensions and resolutions. Photos are stored in Amazon S3 in their original format and resolution. The company uses an Amazon CloudFront distribution to serve the photos The app includes the dimension and resolution of the display as GET parameters with every request.
A developer needs to implement a solution that optimizes the photos that are served to each device to reduce load time and increase photo quality.
Which solution will meet these requirements MOST cost-effective?

A. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolution
B. Create a dynamic CloudFront origin that automatically maps the request of each device to the corresponding photo variant.
C. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolution
D. Create a Lambda@Edge function to route requests to the corresponding photo vacant by using request headers.
E. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a respons
F. Change the CloudFront TTL cache policy to the maximum value possible.

Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a respons
G: In the same function store a copy of the processed photos on Amazon S3 for subsequent requests.

**Answer:** D

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

**NEW QUESTION 83**
A developer migrated a legacy application to an AWS Lambda function. The function uses a third-party service to pull data with a series of API calls at the end of each month. The function than processes the data to generate the monthly reports. The function has Been working with no issues so far.
The third-party service recently issued a restriction to allow a feed number to API calls each minute and each day. If the API calls exceed the limit tor each minute or each day, then the service will produce errors. The API also provides the minute limit and daily limit in the response header. This restriction might extend the overall process to multiple days because the process is consuming more API calls than the available limit.
What is the MOST operationally efficient way to refactor the server less application to accommodate this change?

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

**Answer:** A

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

**NEW QUESTION 88**
A company is using Amazon RDS as the Backend database for its application. After a recent marketing campaign, a surge of read requests to the database increased the latency of data retrieval from the database.
The company has decided to implement a caching layer in front of the database. The cached content must be encrypted and must be highly available.
Which solution will meet these requirements?

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

**Answer:** C

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

**NEW QUESTION 92**
A developer is creating an AWS Lambda function in VPC mode An Amazon S3 event will invoke the Lambda function when an object is uploaded into an S3 bucket The Lambda function will process the object and produce some analytic results that will be recorded into a file Each processed object will also generate a log entry that will be recorded into a file.
Other Lambda functions. AWS services, and on-premises resources must have access to the result files and log file. Each log entry must also be appended to the same shared log file. The developer needs a solution that can share files and append results into an existing file.
Which solution should the developer use to meet these requirements?

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

**Answer:** A

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

**NEW QUESTION 97**
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

**Answer:** D

**Explanation:**
 The unified Amazon CloudWatch agent can collect both system metrics and log files from Amazon EC2 instances and on-premises servers. By installing and configuring the agent on the EC2 instances, the developer can easily access and analyze the application logs in CloudWatch without logging in to each server individually. This option requires minimum changes to the existing application and does not affect its availability or scalability. References
? Using the CloudWatch Agent
? Collecting Metrics and Logs from Amazon EC2 Instances and On-Premises Servers with the CloudWatch Agent

**NEW QUESTION 102**
A developer is creating an Amazon DynamoDB table by using the AWS CLI The DynamoDB table must use server-side encryption with an AWS owned encryption key
How should the developer create the DynamoDB table to meet these requirements?

A. Create an AWS Key Management Service (AWS KMS) customer managed ke
B. Provide the key's Amazon Resource Name (ARN) in the KMSMasterKeyId parameter during creation of the DynamoDB table
C. Create an AWS Key Management Service (AWS KMS) AWS managed key Provide the key's Amazon Resource Name (ARN) in the KMSMasterKeyId parameter during creation of the DynamoDB table
D. Create an AWS owned key Provide the key's Amazon Resource Name (ARN) in the KMSMasterKeyId parameter during creation of the DynamoDB table.
E. Create the DynamoDB table with the default encryption options

**Answer:** D

**Explanation:**
 When creating an Amazon DynamoDB table using the AWS CLI, server-side encryption with an AWS owned encryption key is enabled by default. Therefore, the developer does not need to create an AWS KMS key or specify the KMSMasterKeyId parameter. Option A and B are incorrect because they suggest creating customer- managed and AWS-managed KMS keys, which are not needed in this scenario. Option C is also incorrect because AWS owned keys are automatically used for server-side encryption by default.

**NEW QUESTION 105**
A developer is building an application that gives users the ability to view bank account from multiple sources in a single dashboard. The developer has automated the process to retrieve API credentials for these sources. The process invokes an AWS Lambda function that is associated with an AWS CloudFormation cotton resource.
The developer wants a solution that will store the API credentials with minimal operational overhead.
When solution will meet these requirements?

A. Add an AWS Secrets Manager GenerateSecretString resource to the CloudFormation templat
B. Set the value to reference new credentials to the Cloudformation resource.
C. Use the AWS SDK ssm PutParameter operation in the Lambda function from the existing, custom resource to store the credentials as a paramete
D. Set the parameter value to reference the new credential
E. Set ma parameter type to SecureString.
F. Add an AWS Systems Manager Parameter Store resource to the CloudFormation templat
G. Set the CloudFormation resource value to reference the new credentials Set the resource NoEcho attribute to true.
H. Use the AWS SDK ssm PutParameter operation in the Lambda function from the existing custom resources to store the credentials as a paramete
I. Set the parameter value to reference the new credential
J. Set the parameter NoEcho attribute to true.

**Answer:** B

**Explanation:**
 The solution that will meet the requirements is to use the AWS SDK ssm PutParameter operation in the Lambda function from the existing custom resource to store the credentials as a parameter. Set the parameter value to reference the new credentials. Set the parameter type to SecureString. This way, the developer can store the API credentials with minimal operational overhead, as AWS Systems Manager Parameter Store provides secure and scalable storage for configuration data. The SecureString parameter type encrypts the parameter value with AWS Key Management Service (AWS KMS). The other options either involve adding additional resources to the CloudFormation template, which increases complexity and cost, or do not encrypt the parameter value, which reduces security.
Reference: Creating Systems Manager parameters

**NEW QUESTION 108**
A company is planning to use AWS CodeDeploy to deploy an application to Amazon Elastic Container Service (Amazon ECS) During the deployment of a new version of the application, the company initially must expose only 10% of live traffic to the new version of the deployed application. Then, after 15 minutes elapse, the company must route all the remaining live traffic to the new version of the deployed application.

Which CodeDeploy predefined configuration will meet these requirements?

A. CodeDeployDefault ECSCanary10Percent15Minutes
B. CodeDeployDefault LambdaCanary10Percent5Minutes
C. CodeDeployDefault LambdaCanary10Percent15Minutes
D. CodeDeployDefault ECSLinear10PercentEvery1 Minutes

**Answer:** A

**Explanation:**
 The predefined configuration "CodeDeployDefault.ECSCanary10Percent15Minutes" is designed for Amazon Elastic Container Service (Amazon ECS) deployments and meets the specified requirements. It will perform a canary deployment, which means it will initially route 10% of live traffic to the new version of the application, and then after 15 minutes elapse, it will automatically route all the remaining live traffic to the new version. This gradual deployment approach allows

the company to verify the health and performance of the new version with a small portion of traffic before fully deploying it to all users.

**NEW QUESTION 113**
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 bucke
C. Mount the S3 bucket folder in the Lambda function Import the library by using the proper folder in the mount point.
D. Load the external library to the Lambda function's /tmp directory during deployment of the Lambda packag
E. Import the library from the /tmp directory.
F. Create an Amazon Elastic File System (Amazon EFS) volum
G. Upload the external library to the EFS volume Mount the EFS volume in the Lambda functio
H. Import the library by using the proper folder in the mount point.

**Answer:** A

**Explanation:**

Create a Lambda layer to store the external library. Configure the Lambda function to use the layer. This will allow the developer to make the external library available to the Lambda execution environment without having to include it in the Lambda package, which will reduce the Lambda package space. Using a Lambda layer is a simple and straightforward solution that requires minimal operational overhead. https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html

**NEW QUESTION 116**
An ecommerce company is using an AWS Lambda function behind Amazon API Gateway

as its application tier. To process orders during checkout, the application calls a POST API from the frontend. The POST API invokes the Lambda function asynchronously. In rare situations, the application has not processed orders. The Lambda application logs show no errors or failures.
What should a developer do to solve this problem?

A. Inspect the frontend logs for API failure
B. Call the POST API manually by using the requests from the log file.
C. Create and inspect the Lambda dead-letter queu
D. Troubleshoot the failed function
E. Reprocess the events.
F. Inspect the Lambda logs in Amazon CloudWatch for possible error
G. Fix the errors.
H. Make sure that caching is disabled for the POST API in API Gateway.

**Answer:** B

**Explanation:**
The solution that will solve this problem is to create and inspect the Lambda dead-letter queue. Troubleshoot the failed functions. Reprocess the events. This way, the developer can identify and fix any issues that caused the Lambda function to fail when invoked asynchronously by API Gateway. The developer can also reprocess any orders that were not processed due to failures. The other options either do not address the root cause of the problem, or do not help recover from failures.
Reference: Asynchronous invocation

**NEW QUESTION 118**
A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the AWS Key Management Service (AWS KMS) key was used and by whom.
Which encryption option will meet these requirements?

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

**Answer:** B

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

**NEW QUESTION 119**
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 bucke
C. Use the S3 Standard storage clas
D. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Flexible Retrieval storage class after 1 year.
E. 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
F. Store the files on an Amazon Elastic File System (Amazon EFS) moun
G. Configure EFS lifecycle management to transition the files to the EFS Standard-Infrequent Access (Standard-IA) storage class after 1 year.

**Answer:** A

**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/

**NEW QUESTION 121**
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 EventBridg
D. Create an EventBridge rule to invoke an AWS Lambda function to process the documents.
E. Update the application to synchronously process the documents directly after the DynamoDB write

**Answer:** B

**Explanation:**
https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and- design-patterns/

**NEW QUESTION 123**
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 code
C. Create API Gateway resources and set the integration type value to AWS_PROXY Deploy the API.
D. 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.
E. Create API Gateway resources and set the integration type value set to HTTP_PROX
F. Add mapping templates and deploy the AP
G. Create an AWS Lambda layer that returns various HTTP status codes Associate the Lambda layer with the API deployment

**Answer:** A

**Explanation:**
The best solution for publishing an API without an integrated backend is to use the MOCK integration type in API Gateway. This allows the developer to return a static response to the client without sending the request to a backend service. The developer can configure the method integration request and integration response to associate a response with an HTTP status code, such as 200 OK or 404 Not Found. The developer can also create an API Gateway stage and deploy the API to make it available to the teams that depend on the application backend. The other solutions are either not feasible or not efficient. Creating an AWS Lambda function, an EC2 application, or an AWS Lambda layer would require additional resources and code to generate the mocked responses and HTTP status codes. These solutions would also incur additional costs and complexity, and would not leverage the built-in functionality of API Gateway. References
? Set up mock integrations for API Gateway REST APIs
? Mock Integration for API Gateway - AWS CloudFormation
? Mocking API Responses with API Gateway
? How to mock API Gateway responses with AWS SAM

**NEW QUESTION 125**
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 tim
B. 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 operatio
C. Schedule a cron job on an Amazon EC2 instance once an hour to start the script.
D. For each item add a new attribute of typ
E. String that has a timestamp that its set to the blog post creation tim

F. 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 operatin
G. Place the script in a container imag
H. Schedule an Amazon Elastic Container Service (Amazon ECS) task on AWS Far gate that invokes the container every 5 minutes.
I. 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 tim
J. Create a global secondary index (GSI) that uses the new attribute as a sort ke
K. 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.
L. For each item add a new attribute of typ
M. Number that has timestamp that is set to 48 hours after the blog pos
N. creation time Configure the DynamoDB table with a TTL that references the new attribute.

**Answer:** D

**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.
References: Time To Live, Managing DynamoDB Time To Live (TTL)


**NEW QUESTION 130**
A developer is optimizing an AWS Lambda function and wants to test the changes in production on a small percentage of all traffic. The Lambda function serves requests to a REST API in Amazon API Gateway. The developer needs to deploy their changes and perform a test in production without changing the API Gateway URL.
Which solution will meet these requirements?

A. Define a function version for the currently deployed production Lambda functio
B. Update the API Gateway endpoint to reference the new Lambda function versio
C. Upload and publish the optimized Lambda function cod
D. On the production API Gateway stage, define a canary release and set the percentage of traffic to direct to the canary releas
E. Update the API Gateway endpoint to use the $LATEST version of the Lambda functio
F. Publish the API to the canary stage.
G. Define a function version for the currently deployed production Lambda functio
H. Update the API Gateway endpoint to reference the new Lambda function versio
I. Upload and publish the optimized Lambda function cod
J. Update the API Gateway endpoint to use the $LATEST version of the Lambda functio
K. Deploy a new API Gateway stage.
L. Define an alias on the $LATEST version of the Lambda functio
M. Update the API Gateway endpoint to reference the new Lambda function alia
N. Upload and publish the optimized Lambda function cod
O. On the production API Gateway stage, define a canary release and set the percentage of traffic to direct to the canary releas
P. Update the API Gateway endpoint to use the SLAT EST version of the Lambda functio
Q. Publish to the canary stage.
R. Define a function version for the currently deployed production Lambda functio
S. Update the API Gateway endpoint to reference the new Lambda function versio
T. Upload and publish the optimized Lambda function cod
. Update the API Gateway endpoint to use the$LATEST version of the Lambda functio
. Deploy the API to the production API Gateway stage.

**Answer:** C

**Explanation:**
? A Lambda alias is a pointer to a specific Lambda function version or another alias1. A Lambda alias allows you to invoke different versions of a function using the same name1. You can also split traffic between two aliases by assigning weights to them1.
? In this scenario, the developer needs to test their changes in production on a small percentage of all traffic without changing the API Gateway URL. To achieve this, the developer can follow these steps:
? By using this solution, the developer can test their changes in production on a small percentage of all traffic without changing the API Gateway URL. The developer can also monitor and compare metrics between the canary and production releases, and promote or disable the canary as needed2.


**NEW QUESTION 132**
A developer needs to migrate an online retail application to AWS to handle an anticipated increase in traffic. The application currently runs on two servers: one server for the web application and another server for the database. The web server renders webpages and manages session state in memory. The database server hosts a MySQL database that contains order details. When traffic to the application is heavy, the memory usage for the web server approaches 100% and the application slows down considerably.
The developer has found that most of the memory increase and performance decrease is related to the load of managing additional user sessions. For the web server migration, the developer will use Amazon EC2 instances with an Auto Scaling group behind an Application Load Balancer.
Which additional set of changes should the developer make to the application to improve the application's performance?

A. Use an EC2 instance to host the MySQL databas
B. Store the session data and the application data in the MySQL database.
C. Use Amazon ElastiCache for Memcached to store and manage the session dat
D. Use an Amazon RDS for MySQL DB instance to store the application data.
E. Use Amazon ElastiCache for Memcached to store and manage the session data and the application data.
F. Use the EC2 instance store to manage the session dat
G. Use an Amazon RDS for MySQL DB instance to store the application data.

**Answer:** B

**Explanation:**
 Using Amazon ElastiCache for Memcached to store and manage the session data will reduce the memory load and improve the performance of the web server. Using Amazon RDS for MySQL DB instance to store the application data will provide a scalable, reliable, and managed database service. Option A is not optimal because it does not address the memory issue of the web server. Option C is not optimal because it does not provide a persistent storage for the application data. Option D is not optimal because it does not provide a high availability and durability for the session data.
References: Amazon ElastiCache, Amazon RDS


**NEW QUESTION 133**
An application uses an Amazon EC2 Auto Scaling group. A developer notices that EC2 instances are taking a long time to become available during scale-out events. The UserData script is taking a long time to run.
The developer must implement a solution to decrease the time that elapses before an EC2 instance becomes available. The solution must make the most recent version of the application available at all times and must apply all available security updates. The solution also must minimize the number of images that are created. The images must be validated.
Which combination of steps should the developer take to meet these requirements? (Choose two.)

A. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install all the patches and agents that are needed to manage and run the applicatio
B. Update the Auto Scaling group launch configuration to use the AMI.
C. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install the latest version of the application and all the patches and agents that are needed to manage and run the applicatio
D. Update the Auto Scaling group launch configuration to use the AMI.
E. Set up AWS CodeDeploy to deploy the most recent version of the application at runtime.
F. Set up AWS CodePipeline to deploy the most recent version of the application at runtime.
G. Remove any commands that perform operating system patching from the UserData script.

**Answer:** BE

**Explanation:**
 AWS CloudFormation is a service that enables developers to model and provision AWS resources using templates. The developer can use the following steps to avoid accidental database deletion in the future:
? Set up AWS CodeDeploy to deploy the most recent version of the application at
runtime. This will ensure that the application code is always up to date and does not depend on the AMI.
? Remove any commands that perform operating system patching from the
UserData script. This will reduce the time that the UserData script takes to run and speed up the instance launch process.
References:
? [What Is AWS CloudFormation? - AWS CloudFormation]
? [What Is AWS CodeDeploy? - AWS CodeDeploy]
? [Running Commands on Your Linux Instance at Launch - Amazon Elastic Compute Cloud]


**NEW QUESTION 137**
A developer is writing a serverless application that requires an AWS Lambda function to be invoked every 10 minutes.
What is an automated and serverless way to invoke the function?

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

**Answer:** C

**Explanation:**
 The solution that will meet the requirements is to create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function. This way, the developer can use an automated and serverless way to invoke the function every 10 minutes. The developer can also use a cron expression or a rate expression

to specify the schedule for the rule. The other options either involve using an Amazon EC2 instance, which is not serverless, or using environment variables or query parameters, which do not trigger the function.
Reference: Schedule AWS Lambda functions using EventBridge

## NEW QUESTION 139

A developer is building a serverless application that is based on AWS Lambda. 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

**Answer:** B

**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]

## NEW QUESTION 143

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 service
B. Update the task role policy to allow access to me X -Ray API.
C. Deploy AWS X-Ray as a daemon set to the Fargate cluste
D. Update the service role

policy to allow access to the X-Ray API.
E. Instrument the application by using the AWS X-Ray SD
F. Update the application to use the Put-XrayTrace API call to communicate with the X-Ray API.
G. Instrument the application by using the AWS X-Ray SD
H. Update the application to communicate with the X-Ray daemon.
I. Instrument the ECS task to send the stout and spider- output to Amazon CloudWatch Log
J. Update the task role policy to allow the cloudwatch Putlogs action.

**Answer:** AE

**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

**NEW QUESTION 144**
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

**Answer:** B

**Explanation:**
https://docs.aws.amazon.com/step-functions/latest/dg/welcome.html

**NEW QUESTION 145**

......

# Thank You for Trying Our Product

* 100% Pass or Money Back

    All our products come with a 90-day Money Back Guarantee.

* One year free update

    You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

    We currently serve more than 30,000,000 customers.

* Shop Securely

    All transactions are protected by VeriSign!

**100% Pass Your AWS-Certified-Developer-Associate Exam with Our Prep Materials Via below:**

https://www.certleader.com/AWS-Certified-Developer-Associate-dumps.html