Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

NEW QUESTION 1

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 2

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 3

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

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

Answer: C

Explanation:

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

NEW QUESTION 4

A 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 5

An application that runs on AWS receives messages from an Amazon Simple Queue Service (Amazon SQS) queue and processes the messages in batches. The   application sends the data to another SQS queue to be consumed by another legacy application. The legacy system can take up to 5 minutes to process some transaction data.

A developer wants to ensure that there are no out-of-order updates in the legacy system. The developer cannot alter the behavior of the legacy system.

Which solution will meet these requirements?

A. Use an SQS FIFO queue

B. Configure the visibility timeout value.

C. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ

D. Configure the DelaySeconds values.

E. Use an SQS standard queue with a SendMessageBatchRequestEntry data typ

F. Configure the visibility timeout value.

G. Use an SQS FIFO queue

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 6

A developer is creating an AWS Lambda function that needs credentials to connect to an Amazon RDS for MySQL database. An Amazon S3 bucket currently stores the credentials. The developer needs to improve the existing solution by implementing credential rotation and secure storage. The developer also needs to provide integration with the Lambda function.   Which solution should the developer use to store and retrieve the credentials with the LEAST management overhead?

A. Store the credentials in AWS Systems Manager Parameter Stor

B. Select the database that the parameter will acces

C. Use the default AWS Key Management Service (AWS KMS) key to encrypt the paramete

D. Enable automatic rotation for the paramete

E. Use the parameter from Parameter Store on the Lambda function to connect to the database.

F. Encrypt the credentials with the default AWS Key Management Service (AWS KMS) ke

G. Store the credentials as environment variables for the Lambda functio

H. Create a second Lambda function to generate new credentials and to rotate the credentials by updating the environment variables of the first Lambda functio

I. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedul

J. Update the database to use the new credential

K. On the first Lambda function, retrieve the credentials from the environment variable

L. Decrypt the credentials by using AWS KMS, Connect to the database.

M. Store the credentials in AWS Secrets Manage

N. Set the secret type to Credentials for Amazon RDS databas

O. Select the database that the secret will acces
P. Use the default AWS Key Management Service (AWS KMS) key to encrypt the secre
Q. Enable automatic rotation for the secre
R. Use the secret from Secrets Manager on the Lambda function to connect to the database.
S. Encrypt the credentials by using AWS Key Management Service (AWS KMS). Store the credentials in an Amazon DynamoDB tabl
T. Create a second Lambda function to rotate the credential
. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedul
. Update the DynamoDB tabl
. Update the database to use the generated credential
. Retrieve the credentials from DynamoDB with the first Lambda functio
. Connect to the database.

Answer: C

Explanation:
 AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. Secrets Manager enables you to store, retrieve, and rotate secrets such as database credentials, API keys, and passwords. Secrets Manager supports a secret type for RDS databases, which allows you to select an existing RDS database instance and generate credentials for it. Secrets Manager encrypts the secret using AWS Key Management Service (AWS KMS) keys and enables automatic rotation of the secret at a specified interval. A Lambda function can use the AWS SDK or CLI to retrieve the secret from Secrets Manager and use it to connect to the database. Reference: Rotating your AWS Secrets Manager secrets

NEW QUESTION 7
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
framework for the other developers on the tea
F. Check the response Document how to run the unit testing.
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 8
An online sales company is developing a serverless application that runs on AWS. The application uses an AWS Lambda function that calculates order success rates and stores the data in an Amazon DynamoDB table. A developer wants an efficient way to invoke the Lambda function every 15 minutes.
Which solution will meet this requirement with the LEAST development effort?

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

Answer: A

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

NEW QUESTION 9

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 10
A developer is creating a new REST API by using Amazon API Gateway and AWS Lambda. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.
The developer wants to make the REST API available for testing by using API Gateway   locally.
Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

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

Answer: D

Explanation:
? The sam local start-api subcommand allows you to run your serverless application locally for quick development and testing1. It creates a local HTTP server that acts as a proxy for API Gateway and invokes your Lambda functions based on the AWS SAM template1. You can use the sam local start-api subcommand to test your REST API locally by sending HTTP requests to the local endpoint1.

NEW QUESTION 10
A developer is creating an AWS CloudFormation template to deploy Amazon EC2 instances across multiple AWS accounts. The developer must choose the EC2 instances from a list of approved instance types.
How can the developer incorporate the list of approved instance types in the CloudFormation template?

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

Answer: D

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

NEW QUESTION 12
A developer has an application that makes batch requests directly to Amazon DynamoDB by using the BatchGetItem low-level API operation. The responses frequently return values in the UnprocessedKeys element.
Which actions should the developer take to increase the resiliency of the application when the batch response includes values in UnprocessedKeys? (Choose two.)

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

Answer: BC

Explanation:
The UnprocessedKeys element indicates that the BatchGetItem operation did not process all of the requested items in the current response. This can happen if the
                              response size limit is exceeded or if the table's provisioned throughput is exceeded. To handle this situation, the developer should retry

the batch operation with exponential backoff and randomized delay to avoid throttling errors and reduce the load on the table. The developer should also use an AWS SDK to make the requests, as the SDKs automatically retry requests that return UnprocessedKeys.
References:
? [BatchGetItem - Amazon DynamoDB]
? [Working with Queries and Scans - Amazon DynamoDB]
? [Best Practices for Handling DynamoDB Throttling Errors]

NEW QUESTION 14
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 19
A developer is creating an application that will give users the ability to store photos from their cellphones in the cloud. The application needs to support tens of thousands of users. The application uses an Amazon API Gateway REST API that is integrated with AWS Lambda functions to process the photos. The application stores details about the photos in Amazon DynamoDB.
Users need to create an account to access the application. In the application, users must be able to upload photos and retrieve previously uploaded photos. The photos will range in size from 300 KB to 5 MB.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Cognito user pools to manage user account
B. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the AP
C. Use the Lambda function to store the photos and details in the DynamoDB tabl
D. Retrieve previously uploaded photos directly from the DynamoDB table.
E. Use Amazon Cognito user pools to manage user account
F. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the AP
G. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB tabl
H. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
I. Create an IAM user for each user of the application during the sign-up proces
J. Use IAM authentication to access the API Gateway AP
K. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB tabl
L. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
M. Create a users table in DynamoD
N. Use the table to manage user account
O. Create a Lambda authorizer that validates user credentials against the users tabl
P. Integrate the Lambda authorizer with API Gateway to control access to the AP
Q. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as par of the photo details in the DynamoDB tabl
R. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.

Answer: B

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

NEW QUESTION 23
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 28
A company runs an application on AWS The application stores data in an Amazon DynamoDB table Some queries are taking a long time to run These slow queries involve an attribute that is not the table's partition key or sort key
The amount of data that the application stores in the DynamoDB table is expected to increase significantly. A developer must increase the performance of the queries.
Which solution will meet these requirements'?

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

Answer: B

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

NEW QUESTION 32
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 CllCD 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 34

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 39

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 42

A company wants to automate part of its deployment process. A developer needs to automate the process of checking for and deleting unused resources that supported previously deployed stacks but that are no longer used.

The company has a central application that uses the AWS Cloud Development Kit (AWS CDK) to manage all deployment stacks. The stacks are spread out across multiple accounts. The developer's solution must integrate as seamlessly as possible within the current deployment process.

Which solution will meet these requirements with the LEAST amount of configuration?

A. In the central AWS CDK application, write a handler function in the code that uses AWS SDK calls to check for and delete unused resource

B. Create an AWS CloudPormation template from a JSON fil

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

Answer: B

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

NEW QUESTION 43
A developer at a company needs to create a small application mat makes the same API call once each flay at a designated time. The company does not have infrastructure in the AWS Cloud yet, but the company wants to implement this functionality on AWS.
Which solution meets these requirements in the MOST operationally efficient manner?

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

Answer: C

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

NEW QUESTION 46
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 47
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 49
A company is migrating its PostgreSQL database into the AWS Cloud. The company wants to use a database that will secure and regularly rotate database credentials. The company wants a solution that does not require additional programming overhead.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 51
A company is creating an application that processes csv files from Amazon S3 A developer has created an S3 bucket The developer has also created an AWS Lambda function to process the csv files from the S3 bucket
Which combination of steps will invoke the Lambda function when a csv file is uploaded to Amazon S3? (Select TWO.)

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

Answer: AC

Explanation:
To invoke a Lambda function when a csv file is uploaded to Amazon S3, you can use Amazon EventBridge to create a rule that matches the S3 object created event. Then, you can add a trigger to the existing Lambda function and set the trigger type to EventBridge. This way, the Lambda function will be invoked whenever a new csv file is added to the S3 bucket. References
? Tutorial: Using an Amazon S3 trigger to invoke a Lambda function
? How to trigger my Lambda Function once the file is uploaded to s3 bucket
? Lambda Function to be invoked or triggered by S3(csv file upload …

NEW QUESTION 56
A company has an analytics application that uses an AWS Lambda function to process transaction data asynchronously A developer notices that asynchronous invocations of the Lambda function sometimes fail When failed Lambda function invocations occur, the developer wants to invoke a second Lambda function to handle errors and log details.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 57
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 60
A developer uses AWS CloudFormation to deploy an Amazon API Gateway API and an AWS Step Functions state machine The state machine must reference the API Gateway API after the CloudFormation template is deployed The developer needs a solution that uses the state machine to reference the API Gateway endpoint.
Which solution will meet these requirements MOST cost-effectively?

A. Configure the CloudFormation template to reference the API endpoint in the DefinitionSubstitutions property for the AWS StepFunctions StateMachme resource.
B. Configure the CloudFormation template to store the API endpoint in an environment variable for the AWS::StepFunctions::StateMachine resourc Configure the state machine to reference the environment variable
C. Configure the CloudFormation template to store the API endpoint in a standard AWS: SecretsManager Secret resource Configure the state machine to reference the resource
D. Configure the CloudFormation template to store the API endpoint in a standard AWS::AppConfig;:ConfigurationProfile resource Configure the state machine to reference the resource.

Answer: A

Explanation:

The most cost-effective solution is to use the DefinitionSubstitutions property of the AWS::StepFunctions::StateMachine resource to inject the API endpoint as a variable in the state machine definition. This way, the developer can use the intrinsic function
Fn::GetAtt to get the API endpoint from the AWS::ApiGateway::RestApi resource, and pass it to the state machine without creating any additional resources or environment variables. The other solutions involve creating and managing extra resources, such as Secrets Manager secrets or AppConfig configuration profiles, which incur additional costs and complexity. References
? AWS::StepFunctions::StateMachine - AWS CloudFormation
? Call API Gateway with Step Functions - AWS Step Functions
? amazon-web-services aws-api-gateway terraform aws-step-functions

NEW QUESTION 64
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.
E. 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 68
A developer is working on an ecommerce website The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available
How can the developer update the application to meet these requirements with MINIMUM changes?

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

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 73
An application is using Amazon Cognito user pools and identity pools for secure access. A developer wants to integrate the user-specific file upload and download features in the application with Amazon S3. The developer must ensure that the files are saved and retrieved in a secure manner and that users can access only their own files. The file sizes range from 3 KB to 300 MB.
Which option will meet these requirements with the HIGHEST level of security?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/cognito/latest/developerguide/amazon-cognito-integrating-user-pools-with-identity-pools.html

NEW QUESTION 75
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment, the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

A.

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

Answer: D

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

NEW QUESTION 80
An ecommerce application is running behind an Application Load Balancer. A developer observes some unexpected load on the application during non-peak hours. The developer wants to analyze patterns for the client IP addresses that use the application. Which HTTP header should the developer use for this analysis?

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

Answer: C

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

NEW QUESTION 84
A developer is creating a template that uses AWS CloudFormation to deploy an application. The application is serverless and uses Amazon API Gateway, Amazon DynamoDB, and AWS Lambda.
Which AWS service or tool should the developer use to define serverless resources in YAML?

A. CloudFormation serverless intrinsic functions
B. AWS Elastic Beanstalk
C. AWS Serverless Application Model (AWS SAM)
D. AWS Cloud Development Kit (AWS CDK)

Answer: C

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

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

NEW QUESTION 89
A developer is incorporating AWS X-Ray into an application that handles personal
identifiable information (PII). The application is hosted on Amazon EC2 instances. The application trace messages include encrypted PII and go to Amazon CloudWatch. The developer needs to ensure that no PII goes outside of the EC2 instances.
Which solution will meet these requirements?

A. Manually instrument the X-Ray SDK in the application code.
B. Use the X-Ray auto-instrumentation agent.
C. Use Amazon Macie to detect and hide PI
D. Call the X-Ray API from AWS Lambda.
E. Use AWS Distro for Open Telemetry.

Answer: A

Explanation:
This solution will meet the requirements by allowing the developer to control what data is sent to X-Ray and CloudWatch from the application code. The developer can filter out any PII from the trace messages before sending them to X-Ray and CloudWatch, ensuring that no PII goes outside of the EC2 instances. Option B is not optimal because it will automatically instrument all incoming and outgoing requests from the application, which may include PII in the trace messages. Option C is not optimal because it will require additional services and costs to use Amazon Macie and AWS Lambda, which may not be able to detect and hide all PII from the trace messages. Option D is not optimal because it will use Open Telemetry instead of X-Ray, which may not be compatible with CloudWatch and other AWS services.
References: [AWS X-Ray SDKs]

NEW QUESTION 93
A developer needs to perform geographic load testing of an API. The developer must deploy resources to multiple AWS Regions to support the load testing of the API.
How can the developer meet these requirements without additional application code?

A. Create and deploy an AWS Lambda function in each desired Regio
B. Configure the Lambda function to create a stack from an AWS CloudFormation template in that Region when the function is invoked.
C: Create an AWS CloudFormation template that defines the load test resource
D: Use the AWS CLI create-stack-set command to create a stack set in the desired Regions.
E. Create an AWS Systems Manager document that defines the resource
F. Use the document to create the resources in the desired Regions.
G. Create an AWS CloudFormation template that defines the load test resource
H. Use the AWS CLI deploy command to create a stack from the template in each Region.

Answer: B

Explanation:
 AWS CloudFormation is a service that allows developers to model and provision AWS resources using templates. A CloudFormation template can define the load test resources, such as EC2 instances, load balancers, and Auto Scaling groups. A CloudFormation stack set is a collection of stacks that can be created and managed from a single template in multiple Regions and accounts. The AWS CLI create-stack-set command can be used to create a stack set from a template and specify the Regions where the stacks should be created. Reference: Working with AWS CloudFormation stack sets

NEW QUESTION 97
A company has an application that uses Amazon Cognito user pools as an identity provider. The company must secure access to user records. The company has set up multi-factor authentication (MFA). The company also wants to send a login activity notification by email every time a user logs in.
What is the MOST operationally efficient solution that meets this requirement?

A. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notificatio
B. Add an Amazon API Gateway API to invoke the functio
C. Call the API from the client side when login confirmation is received.
D. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notificatio
E. Add an Amazon Cognito post authentication Lambda trigger for the function.
F. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notificatio
G. Create an Amazon CloudWatch Logs log subscription filter to invoke the function based on the login status.
H. Configure Amazon Cognito to stream all logs to Amazon Kinesis Data Firehos
I. Create an AWS Lambda function to process the streamed logs and to send the email notification based on the login status of each user.

Answer: B

Explanation:
 Amazon Cognito user pools support Lambda triggers, which are custom functions that can be executed at various stages of the user pool workflow. A post authentication Lambda trigger can be used to perform custom actions after a user is authenticated, such as sending an email notification. Amazon SES is a cloud-based email sending service that can be used to send transactional or marketing emails. A Lambda function can use the Amazon SES API to send an email to the user's email address after the user logs in successfully. Reference: Post authentication Lambda trigger

NEW QUESTION 100
A developer is building an application that uses AWS API Gateway APIs. AWS Lambda function, and AWS Dynamic DB tables. The developer uses the AWS Serverless Application Model (AWS SAM) to build and run serverless applications on AWS. Each time the developer pushes of changes for only to the Lambda functions, all the artifacts in the application are rebuilt.
The developer wants to implement AWS SAM Accelerate by running a command to only redeploy the Lambda functions that have changed.
Which command will meet these requirements?

A. sam deploy -force-upload
B. sam deploy -no-execute-changeset
C. sam package
D. sam sync -watch

Answer: D

Explanation:
 The command that will meet the requirements is sam sync -watch. This command enables AWS SAM Accelerate mode, which allows the developer to only redeploy the Lambda functions that have changed. The -watch flag enables file watching, which automatically detects changes in the source code and triggers a redeployment. The other commands either do not enable AWS SAM Accelerate mode, or do not redeploy the Lambda functions automatically.
Reference: AWS SAM Accelerate

NEW QUESTION 105
A company is using Amazon OpenSearch Service to implement an audit monitoring system. A developer needs to create an AWS Cloudformation custom resource that is
associated with an AWS Lambda function to configure the OpenSearch Service domain. The Lambda function must access the OpenSearch Service domain by using Open Search Service internal master user credentials.
What is the MOST secure way to pass these credentials to the Lambdas function?

A. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment variabl
B. Set the No Echo attenuate to true.
C. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and to create a paramete
D. In AWS Systems Manager Parameter Stor

E. Set the No Echo attribute to tru
F. Create an 1AM role that has the ssm GetParameter permissio
G. Assign me role to the Lambda functio
H. Store me parameter name as the Lambda function's environment variabl
I. Resolve the parameter's value at runtime.
J. Use a CloudFormation parameter to pass the master uses credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment varleWe Encrypt the parameters value by using the AWS Key Management Service (AWS KMS) encrypt command.
K. Use CloudFoimalion to create an AWS Secrets Manager Secre
L. Use a CloudFormation dynamic reference to retrieve the secret's value for the OpenSearch Service domain's MasterUserOption
M. Create an 1AM role that has the secrets manage
N. GetSecretvalue permissio
O. Assign the role to the Lambda Function Store the secrets name as the Lambda function's environment variabl
P. Resole the secret's value at runtime.

Answer: D

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

NEW QUESTION 108
A developer needs to deploy an application running on AWS Fargate using Amazon ECS The application has environment variables that must be passed to a container for the application to initialize.
How should the environment variables be passed to the container?

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

Answer: B

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

NEW QUESTION 111
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 113
A developer is designing a serverless application for a game in which users register and log in through a web browser The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API
The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.
Which solution will meet these requirements'?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/cognito/latest/developerguide/signing-up-users-in-your-app.html

NEW QUESTION 118
A developer needs to build an AWS CloudFormation template that self-populates the AWS Region variable that deploys the CloudFormation template
What is the MOST operationally efficient way to determine the Region in which the template is being deployed?

A. Use the AWS:.Region pseudo parameter
B. Require the Region as a CloudFormation parameter
C. Find the Region from the AWS::StackId pseudo parameter by using the Fn::Split intrinsic function
D. Dynamically import the Region by referencing the relevant parameter in AWS Systems Manager Parameter Store

Answer: A

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section- structure.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter- reference.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter- reference.html

NEW QUESTION 122
A company is building a serverless application on AWS. The application uses an AWS Lambda function to process customer orders 24 hours a day, 7 days a week. The Lambda function calls an external vendor's HTTP API to process payments.
During load tests, a developer discovers that the external vendor payment processing API occasionally times out and returns errors. The company expects that some payment processing API calls will return errors.
The company wants the support team to receive notifications in near real time only when
the payment processing external API error rate exceed 5% of the total number of transactions in an hour. Developers need to use an existing Amazon Simple Notification Service (Amazon SNS) topic that is configured to notify the support team.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 127
A developer maintains applications that store several secrets in AWS Secrets Manager. The applications use secrets that have changed over time. The developer needs to identify required secrets that are still in use. The developer does not want to cause any application downtime.
What should the developer do to meet these requirements?

A. Configure an AWS CloudTrail log file delivery to an Amazon S3 bucke

B. Create an Amazon CloudWatch alarm for the GetSecretValu
C. Secrets Manager API operation requests
D. Create a secrets manager-secret-unused AWS Config managed rul
E. Create an Amazon EventBridge rule to Initiate notification when the AWS Config managed rule is met.
F. Deactivate the applications secrets and monitor the applications error logs temporarily.
G. Configure AWS X-Ray for the application
H. Create a sampling rule lo match the

GetSecretValue Secrets Manager API operation requests.

Answer: B

Explanation:
This solution will meet the requirements by using AWS Config to monitor and evaluate whether Secrets Manager secrets are unused or have been deleted, based on specified time periods. The secrets manager-secret-unused managed rule is a predefined rule that checks whether Secrets Manager secrets have been rotated within a specified number of days or have been deleted within a specified number of days after last accessed date. The Amazon EventBridge rule will trigger a notification when the AWS Config managed rule is met, alerting the developer about unused secrets that can be removed without causing application downtime. Option A is not optimal because it will use AWS CloudTrail log file delivery to an Amazon S3 bucket, which will incur additional costs and complexity for storing and analyzing log files that may not contain relevant information about secret usage. 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 use AWS X-Ray to trace secret usage, which will introduce additional overhead and latency for instrumenting and sampling requests that may not be related to secret usage.
References: [AWS Config Managed Rules], [Amazon EventBridge]

NEW QUESTION 129
A company is preparing to migrate an application to the company's first AWS environment Before this migration, a developer is creating a proof-of-concept application to validate a model for building and deploying container-based applications on AWS.
Which combination of steps should the developer take to deploy the containerized proof-of- concept application with the LEAST operational effort? (Select TWO.)

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To deploy a containerized application on AWS with the least operational effort, the developer should package the application into a container image by using the Docker CLI and upload the image to Amazon ECR, which is a fully managed container registry service. Then, the developer should deploy the application to Amazon ECS on AWS Fargate, which is a serverless compute engine for containers that eliminates the need to provision and manage servers or clusters. Amazon ECS will automatically scale, load balance, and monitor the application. References
? How to Deploy Docker Containers | AWS
? Deploy a Web App Using AWS App Runner
? How to Deploy Containerized Apps on AWS Using ECR and Docker

NEW QUESTION 134

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

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

Answer: D

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

NEW QUESTION 139
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Certified-Developer-Associate Practice Exam Features:

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

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

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

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

100% Actual & Verified — Instant Download, Please Click
Order The AWS-Certified-Developer-Associate Practice Test Here