Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

NEW QUESTION 1
A company has a multi-node Windows legacy application that runs on premises. The application uses a network shared folder as a centralized configuration repository to store configuration files in .xml format. The company is migrating the application to Amazon EC2 instances. As part of the migration to AWS, a developer must identify a solution that provides high availability for the repository.
Which solution will meet this requirement MOST cost-effectively?

A. Mount an Amazon Elastic Block Store (Amazon EBS) volume onto one of the EC2 instance
B. Deploy a file system on the EBS volum
C. Use the host operating system to share a folde
D. Update the application code to read and write configuration files from the shared folder.
E. Deploy a micro EC2 instance with an instance store volum
F. Use the host operating system to share a folde
G. Update the application code to read and write configuration files from the shared folder.
H. Create an Amazon S3 bucket to host the repositor
I. Migrate the existing .xml files to the S3 bucke
J. Update the application code to use the AWS SDK to read and write configuration files from Amazon S3.
K. Create an Amazon S3 bucket to host the repositor
L. Migrate the existing .xml files to the S3 bucke
M. Mount the S3 bucket to the EC2 instances as a local volum
N. Update the application code to read and write configuration files from the disk.

Answer: C

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

NEW QUESTION 2
A developer is troubleshooting an Amazon API Gateway API Clients are receiving HTTP 400 response errors when the clients try to access an endpoint of the API. How can the developer determine the cause of these errors?

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

Answer: D

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

NEW QUESTION 3
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 4
A developer has been asked to create an AWS Lambda function that is invoked any time updates are made to items in an Amazon DynamoDB table. The function has been created and appropriate permissions have been added to the Lambda execution role Amazon DynamoDB streams have been enabled for the table, but the function 15 still not being invoked.
Which option would enable DynamoDB table updates to invoke the Lambda function?

A. Change the StreamViewType parameter value to NEW_AND_OLOJMAGES for the DynamoDB table.
B. Configure event source mapping for the Lambda function.
C. Map an Amazon Simple Notification Service (Amazon SNS) topic to the DynamoDB streams.
D. Increase the maximum runtime (timeout) setting of the Lambda function.

Answer: B

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

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

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

Answer: D

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

NEW QUESTION 8
A company has a web application that is hosted on Amazon EC2 instances The EC2 instances are configured to stream logs to Amazon CloudWatch Logs The company needs to receive an Amazon Simple Notification Service (Amazon SNS) notification when the number of application error messages exceeds a defined threshold within a 5-minute period
Which solution will meet these requirements?

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

Answer: D

Explanation:
The best solution is to create a CloudWatch metric filter to match the application error pattern in the log data. This will allow you to create a custom metric that tracks the number of errors in your application. You can then set up a CloudWatch alarm based on this metric and configure it to send an SNS notification when the number of errors exceeds a defined threshold within a 5-minute period. This solution does not require any changes to your application code or installing any additional agents on your EC2 instances. It also leverages the existing integration between CloudWatch and SNS for sending notifications. References
? Create Metric Filters - Amazon CloudWatch Logs
? Creating Amazon CloudWatch Alarms - Amazon CloudWatch
? How to send alert based on log message on CloudWatch - Stack Overflow

NEW QUESTION 9

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 10
A company is building a compute-intensive application that will run on a fleet of Amazon EC2 instances. The application uses attached Amazon Elastic Block Store (Amazon EBS) volumes for storing data. The Amazon EBS volumes will be created at time of initial deployment. The application will process sensitive information. All of the data must be encrypted. The solution should not impact the application's performance.
Which solution will meet these requirements?

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

Answer: A

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

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 11
A developer is working on a Python application that runs on Amazon EC2 instances. The developer wants to enable tracing of application requests to debug performance issues in the code.
Which combination of actions should the developer take to achieve this goal? (Select TWO)

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

Answer: BE

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

NEW QUESTION 16
A developer has an application that makes batch requests directly to Amazon DynamoDB by using the BatchGetItem low-level API operation. The responses frequently return values in the UnprocessedKeys element.
Which actions should the developer take to increase the resiliency of the application when the batch response includes values in UnprocessedKeys? (Choose two.)

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

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 19
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 23
A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

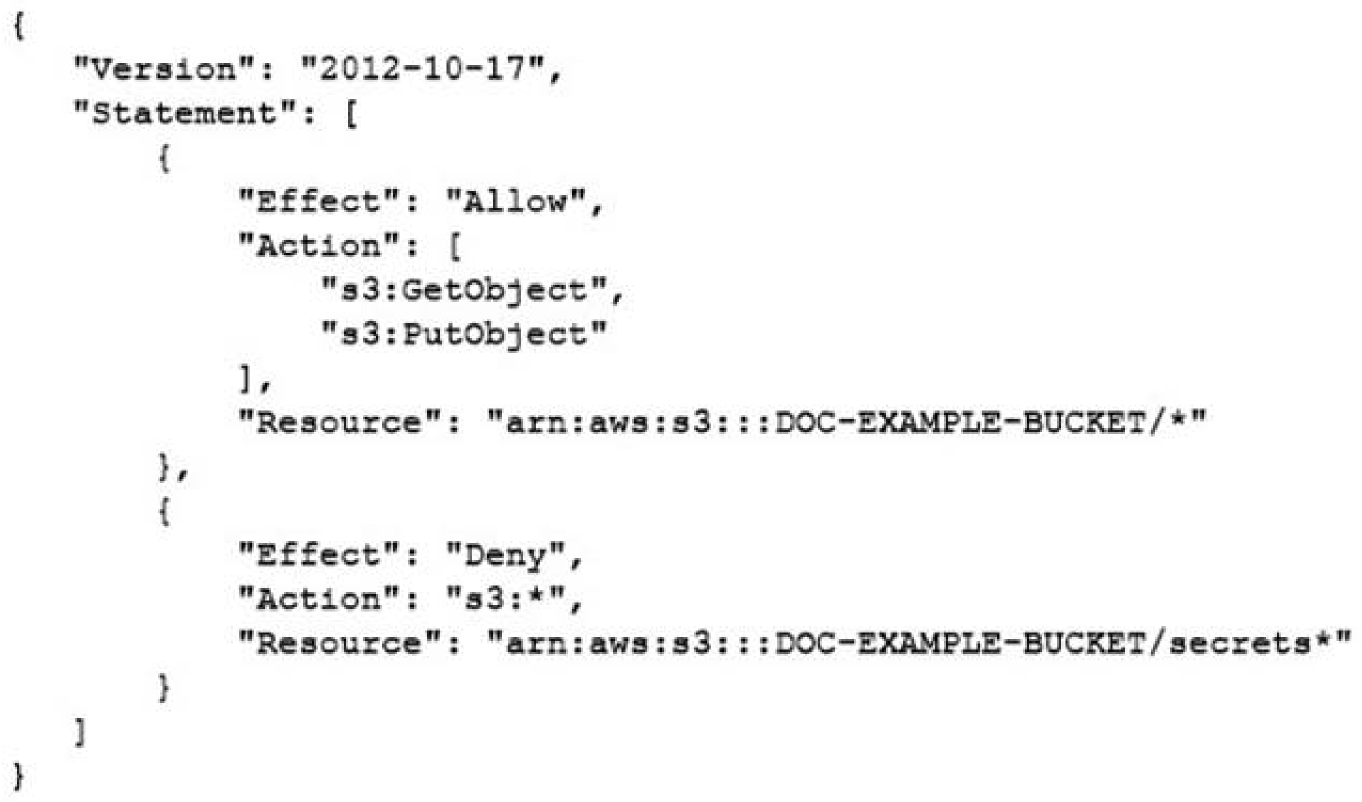

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetObject",
                "s3:PutObject"
            ],
            "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
        },
        {
            "Effect": "Deny",
            "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 24
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 26
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 29
A developer is working on an ecommerce platform that communicates with several third- party payment processing APIs The third-party payment services do not provide a test environment.
The developer needs to validate the ecommerce platform's integration with the third-party payment processing APIs. The developer must test the API integration code without invoking the third-party payment processing APIs.
Which solution will meet these requirements'?

A. Set up an Amazon API Gateway REST API with a gateway response configured for status code 200 Add response templates that contain sample responses captured from the real third-party API.
B. Set up an AWS AppSync GraphQL API with a data source configured for each third- party API Specify an integration type of Mock Configure integration responses by using sample responses captured from the real third-party API.
C. Create an AWS Lambda function for each third-party AP
D. Embed responses captured from the real third-party AP
E. Configure Amazon Route 53 Resolver with an inbound endpoint for each Lambda function's Amazon Resource Name (ARN).
F. Set up an Amazon API Gateway REST API for each third-party API Specify an integration request type of Mock Configure integration responses by using sample responses captured from the real third-party API

Answer: D

Explanation:
Amazon API Gateway can mock responses for testing purposes without requiring any integration backend. This allows the developer to test the API integration code without invoking the third-party payment processing APIs. The developer can configure integration responses by using sample responses captured from the real third- party API. References:
? Mocking Integration Responses in API Gateway
? Set up Mock Integrations for an API in API Gateway

NEW QUESTION 32
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 37
A company wants to share information with a third party. The third party has an HTTP API endpoint that the company can use to share the information. The company has the required API key to access the HTTP API.
The company needs a way to manage the API key by using code. The integration of the API key with the application code cannot affect application performance.
Which solution will meet these requirements MOST securely?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 40
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 AutoPubIishAIias 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 43
A company needs to distribute firmware updates to its customers around the world.
Which service will allow easy and secure control of the access to the downloads at the lowest cost?

A. Use Amazon CloudFront with signed URLs for Amazon S3.
B. Create a dedicated Amazon CloudFront Distribution for each customer.
C. Use Amazon CloudFront with AWS Lambda@Edge.
D. Use Amazon API Gateway and AWS Lambda to control access to an S3 bucket.

Answer: A

Explanation:
This solution allows easy and secure control of access to the downloads at the lowest cost because it uses a content delivery network (CDN) that can cache and distribute firmware updates to customers around the world, and uses a mechanism that can restrict access to specific files or versions. Amazon CloudFront is a CDN that can improve performance, availability, and security of web applications by delivering content from edge locations closer to customers. Amazon S3 is a storage service that can store firmware updates in buckets and objects. Signed URLs are URLs that include additional information, such as an expiration date and time, that give users temporary access to specific objects in S3 buckets. The developer can use CloudFront to serve firmware updates from S3 buckets and use signed URLs to control who can download them and for how long. Creating a dedicated CloudFront distribution for each customer will incur unnecessary costs and complexity. Using Amazon CloudFront with AWS Lambda@Edge will require additional programming overhead to implement custom logic at the edge locations. Using Amazon API Gateway and AWS Lambda to control access to an S3 bucket will also require additional programming overhead and may not provide optimal performance or availability.
Reference: [Serving Private Content through CloudFront], [Using CloudFront with Amazon
S3]

NEW QUESTION 48
A developer is migrating some features from a legacy monolithic application to use AWS Lambda functions instead. The application currently stores data in an Amazon Aurora DB cluster that runs in private subnets in a VPC. The AWS account has one VPC deployed. The Lambda functions and the DB cluster are deployed in the same AWS Region in the same AWS account.
The developer needs to ensure that the Lambda functions can securely access the DB cluster without crossing the public internet.
Which solution will meet these requirements?

A. Configure the DB cluster's public access setting to Yes.
B. Configure an Amazon RDS database proxy for the Lambda functions.
C. Configure a NAT gateway and a security group for the Lambda functions.
D. Configure the VPC, subnets, and a security group for the Lambda functions.

Answer: D

Explanation:
This solution will meet the requirements by allowing the Lambda functions to access the DB cluster securely within the same VPC without crossing the public internet. The developer can configure a VPC endpoint for RDS in a private subnet and assign it to the Lambda functions. The developer can also configure a security group for the Lambda functions that allows inbound traffic from the DB cluster on port 3306 (MySQL). Option A is not optimal because it will expose the DB cluster to public access, which may compromise its security and data integrity. Option B is not optimal because it will introduce additional latency and complexity to use an RDS database proxy for accessing the DB cluster from Lambda functions within the same VPC. Option C is not optimal because it will require additional costs and configuration to use a NAT gateway for accessing resources in private subnets from Lambda functions.
References: [Configuring a Lambda Function to Access Resources in a VPC]

NEW QUESTION 52
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 55
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 60
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 61
A company is implementing an application on Amazon EC2 instances. The application needs to process incoming transactions. When the application detects a transaction that is not valid, the application must send a chat message to the company's support team. To send the message, the application needs to retrieve the access token to authenticate by using the chat API.
A developer needs to implement a solution to store the access token. The access token must be encrypted at rest and in transit. The access token must also be accessible from other AWS accounts.
Which solution will meet these requirements with the LEAST management overhead?

A. Use an AWS Systems Manager Parameter Store SecureString parameter that uses an AWS Key Management Service (AWS KMS) AWS managed key to store the access toke
B. Add a resource-based policy to the parameter to allow access from other account
C. Update the IAM role of the EC2 instances with permissions to access Parameter Stor
D. Retrieve the token from Parameter Store with the decrypt flag enable
E. Use the decrypted access token to send the message to the chat.
F. Encrypt the access token by using an AWS Key Management Service (AWS KMS) customer managed ke
G. Store the access token in an Amazon DynamoDB tabl
H. Update the IAM role of the EC2 instances with permissions to access DynamoDB and AWS KM
I. Retrieve the token from DynamoD
J. Decrypt the token by using AWS KMS on the EC2 instance
K. Use the decrypted access token to send the message to the chat.
L. Use AWS Secrets Manager with an AWS Key Management Service (AWS KMS) customer managed key to store the access toke
M. Add a resource-based policy to the secret to allow access from other account
N. Update the IAM role of the EC2 instances with permissions to access Secrets Manage
O. Retrieve the token from Secrets Manage
P. Use the decrypted access token to send the message to the chat.
Q. Encrypt the access token by using an AWS Key Management Service (AWS KMS) AWS managed ke
R. Store the access token in an Amazon S3 bucke
S. Add a bucket policy to the S3 bucket to allow access from other account
T. Update the IAM role of the EC2 instances with permissions to access Amazon S3 and AWS KM
. Retrieve the token from the S3 bucke
. Decrypt the token by using AWS KMS on the EC2 instance
. Use the decrypted access token to send the massage to the chat.

Answer: C

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/secrets-manager-share-between-accounts/
https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-and- access_examples_cross.html

NEW QUESTION 62
A company built a new application in the AWS Cloud. The company automated the bootstrapping of new resources with an Auto Scaling group by using AWS Cloudf-ormation templates. The bootstrap scripts contain sensitive data.
The company needs a solution that is integrated with CloudFormation to manage the sensitive data in the bootstrap scripts.
Which solution will meet these requirements in the MOST secure way?

A. Put the sensitive data into a CloudFormation paramete
B. Encrypt the CloudFormation templates by using an AWS Key Management Service (AWS KMS) key.
C. Put the sensitive data into an Amazon S3 bucket Update the CloudFormation templates to download the object from Amazon S3 during bootslrap.
D. Put the sensitive data into AWS Systems Manager Parameter Store as a secure string paramete
E. Update the CloudFormation templates to use dynamic references to specify template values.
F. Put the sensitive data into Amazon Elastic File System (Amazon EFS) Enforce EFS encryption after file system creatio
G. Update the CloudFormation templates to retrieve data from Amazon EFS.

Answer: C

Explanation:
This solution meets the requirements in the most secure way because it uses a service that is integrated with CloudFormation to manage sensitive data in encrypted form. AWS Systems Manager Parameter Store provides secure, hierarchical storage for configuration data management and secrets management. You can store sensitive data as secure string parameters, which are encrypted using an AWS Key Management Service (AWS KMS) key of your choice. You can also use dynamic references in your CloudFormation templates to specify template values that are stored in Parameter Store or Secrets Manager without having to include them in your templates. Dynamic references are resolved only during stack creation or update operations, which reduces exposure risks for sensitive data. Putting sensitive data into a CloudFormation parameter will not encrypt them or protect them from unauthorized access. Putting sensitive data into an Amazon S3 bucket or Amazon Elastic File System (Amazon EFS) will require additional configuration and integration with CloudFormation and may not provide fine-grained access control or encryption for sensitive data.
Reference: [What Is AWS Systems Manager Parameter Store?], [Using Dynamic
References to Specify Template Values]

NEW QUESTION 65
A developer is creating an AWS Lambda function that searches for items from an Amazon DynamoDB table that contains customer contact information- The DynamoDB table items have the customer's email_address as the partition key and additional properties such as customer_type, name, and job_tltle.
The Lambda function runs whenever a user types a new character into the customer_type text input The developer wants the search to return partial matches of all the email_address property of a particular customer_type The developer does not want to recreate the DynamoDB table.
What should the developer do to meet these requirements?

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

Answer: A

Explanation:
By adding a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key, the developer can perform a query operation on the GSI using the Begins_with key condition expression with the email_address property. This will return partial matches of all email_address properties of a specific customer_type.

NEW QUESTION 68
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 69
A developer is troubleshooting an application mat uses Amazon DynamoDB in the uswest- 2 Region. The application is deployed to an Amazon EC2 instance. The application requires read-only permissions to a table that is named Cars The EC2 instance has an attached IAM role that contains the following IAM policy.

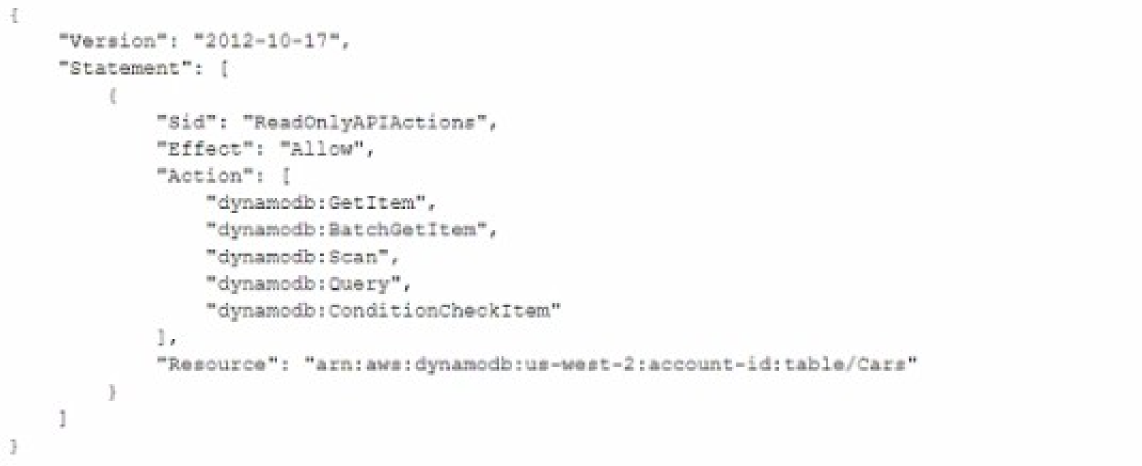

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ReadOnlyAPIActions",
            "Effect": "Allow",
            "Action": [
                "dynamodb:GetItem",
                "dynamodb:BatchGetItem",
                "dynamodb:Scan",
                "dynamodb:Query",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

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

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

Answer: C

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

NEW QUESTION 70
A developer is preparing to begin development of a new version of an application. The previous version of the application is deployed in a production environment. The developer needs to deploy fixes and updates to the current version during the development of the new version of the application. The code for the new version of the application is stored in AWS CodeCommit.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 72
A developer is creating a service that uses an Amazon S3 bucket for image uploads. The service will use an AWS Lambda function to create a thumbnail of each image Each time an image is uploaded the service needs to send an email notification and create the thumbnail The developer needs to configure the image processing and email notifications setup.
Which solution will meet these requirements?

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

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

Answer: A

Explanation:
This solution will allow the developer to receive notifications for each image uploaded to the S3 bucket, and also create a thumbnail using the Lambda function. The SNS topic will serve as a trigger for both the Lambda function and the email notification subscription. When an image is uploaded, S3 will send a notification to the SNS topic, which will trigger the Lambda function to create the thumbnail and also send an email notification to the specified email address.

NEW QUESTION 77
A developer has an application that stores data in an Amazon S3 bucket. The application uses an HTTP API to store and retrieve objects. When the PutObject API operation adds objects to the S3 bucket the developer must encrypt these objects at rest by using server- side encryption with Amazon S3 managed keys (SSE-S3).
Which solution will meet this requirement?

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

Answer: B

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

NEW QUESTION 80
An organization is using Amazon CloudFront to ensure that its users experience low- latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application.
How can these requirements be met? (Select TWO)

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

Answer: BD

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

NEW QUESTION 81
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 86
A developer created an AWS Lambda function that performs a series of operations that involve multiple AWS services. The function's duration time is higher than normal. To determine the cause of the issue, the developer must investigate traffic between the services without changing the function code
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 89
A company needs to set up secure database credentials for all its AWS Cloud resources. The company's resources include Amazon RDS DB instances Amazon DocumentDB clusters and Amazon Aurora DB instances. The company's security policy mandates that database credentials be encrypted at rest and rotated at a regular interval.
Which solution will meet these requirements MOST securely?

A. Set up IAM database authentication for token-based acces
B. Generate user tokens to provide centralized access to RDS DB instance
C. Amazon DocumentDB clusters and Aurora DB instances.
D. Create parameters for the database credentials in AWS Systems Manager Parameter Store Set the Type parameter to Secure Stin
E. Set up automatic rotation on the parameters.
F. Store the database access credentials as an encrypted Amazon S3 object in an S3 bucket Block all public access on the S3 bucke
automatic rotation on the encryption key.
G. Use S3 server-side encryption to set up
H. Create an AWS Lambda function by using the SecretsManagerRotationTemplate template in the AWS Secrets Manager consol
I. Create secrets for the database credentials in Secrets Manager Set up secrets rotation on a schedule.

Answer: D

Explanation:
This solution will meet the requirements by using AWS Secrets Manager, which is a service that helps protect secrets such as database credentials by encrypting them with AWS Key Management Service (AWS KMS) and enabling automatic rotation of secrets. The developer can create an AWS Lambda function by using the SecretsManagerRotationTemplate template in the AWS Secrets Manager console, which provides a sample code for rotating secrets for RDS DB instances, Amazon DocumentDB clusters, and Amazon Aurora DB instances. The developer can also create secrets for the database credentials in Secrets Manager, which encrypts them at rest and provides secure access to them. The developer can set up secrets rotation on a schedule, which changes the database credentials periodically according to a specified interval or event. Option A is not optimal because it will set up IAM database authentication for token-based access, which may not be compatible with all database engines and may require additional configuration and management of IAM roles or users. Option B is not optimal because it will create parameters for the database credentials in AWS Systems Manager Parameter Store, which does not support automatic rotation of secrets. Option C is not optimal because it will store the database access credentials as an encrypted Amazon S3 object in an S3 bucket, which may introduce additional costs and complexity for accessing and securing the data.
References: [AWS Secrets Manager], [Rotating Your AWS Secrets Manager Secrets]

NEW QUESTION 93
A company has an ecommerce application. To track product reviews, the company's development team uses an Amazon DynamoDB table.
Every record includes the following
• A Review ID a 16-digrt universally unique identifier (UUID)
• A Product ID and User ID 16 digit UUIDs that reference other tables
• A Product Rating on a scale of 1-5
• An optional comment from the user
The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product.
Which index will provide the FASTEST response for this query"?

A. A global secondary index (GSI) with Product ID as the partition key and Product Rating as the sort key
B. A global secondary index (GSI) with Product ID as the partition key and Review ID as the sort key
C. A local secondary index (LSI) with Product ID as the partition key and Product Rating as the sort key
D. A local secondary index (LSI) with Review ID as the partition key and Product ID as the sort key

Answer: A

Explanation:
This solution allows the fastest response for the query because it enables the query to use a single partition key value (the Product ID) and a range of sort key values (the Product Rating) to find the matching items. A global secondary index (GSI) is an index that has a partition key and an optional sort key that are different from those on the base table. A GSI can be created at any time and can be queried or scanned independently of the base table. A local secondary index (LSI) is an index that has the same partition key as the base table, but a different sort key. An LSI can only be created when the base table is created and must be queried together with the base table partition key. Using a GSI with Product ID as the partition key and Review ID as the sort key will not allow the query to use a range of sort key values to find the highest ratings. Using an LSI with Product ID as the partition key and Product Rating as the sort key will not work because Product ID is

not the partition key of the base table. Using an LSI with Review ID as the partition key and Product ID as the sort key will not allow the query to use a single partition key value to find the matching items.
Reference: [Global Secondary Indexes], [Querying]

NEW QUESTION 94
A company is running a custom application on a set of on-premises Linux servers that are accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled on the API test stage.
How can a developer enable X-Ray tracing on the on-premises servers with the LEAST amount of configuration?

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

Answer: B

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

NEW QUESTION 97
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 101
A company’s website runs on an Amazon EC2 instance and uses Auto Scaling to scale the environment during peak times. Website users across the world ate experiencing high latency flue lo sialic content on theEC2 instance. even during non-peak hours.
When companion of steps mill resolves the latency issue? (Select TWO)

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

Answer: DE

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

NEW QUESTION 105
A company runs a batch processing application by using AWS Lambda functions and Amazon API Gateway APIs with deployment stages for development, user acceptance testing and production A development team needs to configure the APIs in the deployment stages to connect to third-party service endpoints.
Which solution will meet this requirement?

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

Answer: B

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

NEW QUESTION 106
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 110
A developer is planning to migrate on-premises company data to Amazon S3. The data must be encrypted, and the encryption Keys must support automate annual rotation. The company must use AWS Key Management Service (AWS KMS) to encrypt the data.
When type of keys should the developer use to meet these requirements?

A. Amazon S3 managed keys
B. Symmetric customer managed keys with key material that is generated by AWS
C. Asymmetric customer managed keys with key material that generated by AWS
D. Symmetric customer managed keys with imported key material

Answer: B

Explanation:
The type of keys that the developer should use to meet the requirements is symmetric customer managed keys with key material that is generated by AWS. This way, the developer can use AWS Key Management Service (AWS KMS) to encrypt the data with a symmetric key that is managed by the developer. The developer can also enable automatic annual rotation for the key, which creates new key material for the key every year. The other options either involve using Amazon S3 managed keys, which do not support automatic annual rotation, or using asymmetric keys or imported key material, which are not supported by S3 encryption.
Reference: Using AWS KMS keys to encrypt S3 objects

NEW QUESTION 112
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
F. Update the Lambda function code to download the log file, append the log entries, and upload the modified log file to Amazon EBS
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 115
A company has built an AWS Lambda function to convert large image files into output files that can be used in a third-party viewer application The company recently added a new module to the function to improve the output of the generated files However, the new module has increased the bundle size and has increased the time that is needed to deploy changes to the function code.
How can a developer increase the speed of the Lambda function deployment?

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

Answer: B

Explanation:
Using Lambda layers is a way to reduce the size of the deployment package and speed up the deployment process. Lambda layers are reusable components that can contain libraries, custom runtimes, or other dependencies. By using layers, the developer can separate the core function logic from the dependencies, and avoid uploading them every time the function code changes. Layers can also be shared across multiple functions or accounts, which can improve consistency and maintainability. References
? Working with AWS Lambda layers
? AWS Lambda Layers Best Practices
? Best practices for working with AWS Lambda functions

NEW QUESTION 116
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 120
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 KMSMasterKeyld 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 KMSMasterKeyld parameter during creation of the DynamoDB table
D. Create an AWS owned key Provide the key's Amazon Resource Name (ARN) in the KMSMasterKeyld 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 124
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 128

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

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

Answer: D

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

NEW QUESTION 131
A developer is troubleshooting an application in an integration environment. In the application, an Amazon Simple Queue Service (Amazon SQS) queue consumes messages and then an AWS Lambda function processes the messages. The Lambda function transforms the messages and makes an API call to a third-party service.
There has been an increase in application usage. The third-party API frequently returns an HTTP 429 Too Many Requests error message. The error message prevents a significant number of messages from being processed successfully.
How can the developer resolve this issue?

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

Answer: D

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

NEW QUESTION 132

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 137

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 139
A data visualization company wants to strengthen the security of its core applications The applications are deployed on AWS across its development staging, pre-production, and production environments. The company needs to encrypt all of its stored sensitive credentials The sensitive credentials need to be automatically rotated Aversion of the sensitive credentials need to be stored for each environment
Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: D

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

NEW QUESTION 141
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 146

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

A. Export the existing API to an OpenAPI fil
B. Create a new AP
Modify the new API to add request validatio
C. Import the OpenAPI file.
D. Perform the test
E. Modify the existing API to add request validatio
F. Deploy the existing API to production.
G. Modify the existing API to add request validatio
H. Deploy the updated API to a new API Gateway stag
I. Perform the test
J. Deploy the updated API to the API Gateway production stage.
K. Create a new AP
L. Add the necessary resources and methods, including new request validatio
M. Perform the test
N. Modify the existing API to add request validatio
O. Deploy the existing API to production.
P. Clone the existing AP
Q. Modify the new API to add request validatio
R. Perform the test
S. Modify the existing API to add request validatio
T. Deploy the existing API to production.

Answer: B

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

NEW QUESTION 151
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 154
A development team wants to build a continuous integration/continuous delivery (CI/CD) pipeline. The team is using AWS CodePipeline to automate the code build and deployment. The team wants to store the program code to prepare for the CI/CD pipeline.
Which AWS service should the team use to store the program code?

A. AWS CodeDeploy
B. AWS CodeArtifact
C. AWS CodeCommit
Amazon CodeGuru
D.

Answer: C

Explanation:
AWS CodeCommit is a service that provides fully managed source control for hosting secure and scalable private Git repositories. The development team can use CodeCommit to store the program code and prepare for the CI/CD pipeline. CodeCommit integrates with other AWS services such as CodePipeline, CodeBuild, and CodeDeploy to automate the code build and deployment process.
References:
? [What Is AWS CodeCommit? - AWS CodeCommit]
? [AWS CodePipeline - AWS CodeCommit]

NEW QUESTION 157

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 161
A developer is deploying a new application to Amazon Elastic Container Service (Amazon ECS). The developer needs to securely store and retrieve different types of variables. These variables include authentication information for a remote API, the URL for the API, and credentials. The authentication information and API URL must be available to all current and future deployed versions of the application across development, testing, and production environments.
How should the developer retrieve the variables with the FEWEST application changes?

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

Answer: A

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

NEW QUESTION 162
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 166
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 ltem 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 168
......