Exam Questions DVA-C02

DVA-C02

https://www.2passeasy.com/dumps/DVA-C02/

NEW QUESTION 1

A developer is creating an application that includes an Amazon API Gateway REST API in the us-east-2 Region. The developer wants to use Amazon CloudFront and a custom domain name for the API. The developer has acquired an SSL/TLS certificate for the domain from a third-party provider.
How should the developer configure the custom domain for the application?

A. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the AP
B. Create a DNS A record for the custom domain.
C. Import the SSL/TLS certificate into CloudFron
D. Create a DNS CNAME record for the custom domain.
E. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the AP
F. Create a DNS CNAME record for the custom domain.
G. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the us-east-1 Regio
H. Create a DNS CNAME record for the custom domain.

Answer: D

Explanation:
Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. Amazon CloudFront is a content delivery network (CDN) service that can improve the performance and security of web applications. The developer can use CloudFront and a custom domain name for the API Gateway REST API. To do so, the developer needs to import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the us-east-1 Region. This is because CloudFront requires certificates from ACM to be in this Region. The developer also needs to create a DNS CNAME record for the custom domain that points to the CloudFront distribution.
References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [What Is Amazon CloudFront? - Amazon CloudFront]
? [Custom Domain Names for APIs - Amazon API Gateway]

NEW QUESTION 2

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 3

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 4
A developer needs to deploy an application running on AWS Fargate using Amazon ECS The application has environment variables that must be passed to a container for the application to initialize.
How should the environment variables be passed to the container?

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

Answer: B

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

NEW QUESTION 5
A developer is testing a RESTful application that is deployed by using Amazon API Gateway and AWS Lambda When the developer tests the user login by using credentials that are not valid, the developer receives an HTTP 405 METHOD_NOT_ALLOWED error The developer has verified that the test is sending the correct request for the resource
Which HTTP error should the application return in response to the request?

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

Answer: A

Explanation:
The HTTP 401 error indicates that the request has not been applied because it lacks valid authentication credentials for the target resource. This is the appropriate error code to return when the user login fails due to invalid credentials. The HTTP 405 error means that the method specified in the request is not allowed for the resource identified by the request URI, which is not the case here. The other error codes are not relevant to the authentication failure scenario.
References
? HTTP Status Codes
? AWS Lambda Function Errors in API Gateway

NEW QUESTION 6
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 7
A company has an application that runs across multiple AWS Regions. The application is experiencing performance issues at irregular intervals. A developer must use AWS X-Ray to implement distributed tracing for the application to troubleshoot the root cause of the performance issues.
What should the developer do to meet this requirement?

A. Use the X-Ray console to add annotations for AWS services and user-defined services
B. Use Region annotation that X-Ray adds automatically for AWS services Add Region annotation for user-defined services
C. Use the X-Ray daemon to add annotations for AWS services and user-defined services
D. Use Region annotation that X-Ray adds automatically for user-defined services Configure X-Ray to add Region annotation for AWS services

Answer: B

Explanation:
AWS X-Ray automatically adds Region annotation for AWS services that are integrated with X-Ray. This annotation indicates the AWS Region where the service is running. The developer can use this annotation to filter and group traces by Region and identify any performance issues related to cross-Region calls. The developer can also add Region annotation for user-defined services by using the X-Ray SDK. This option enables the developer to implement distributed tracing for the application that runs across multiple AWS Regions. References
? AWS X-Ray Annotations
? AWS X-Ray Concepts

NEW QUESTION 8
A developer at a company recently created a serverless application to process and show data from business reports. The application's user interface (UI) allows users to select and start processing the files. The Ul displays a message when the result is available to view. The application uses AWS Step Functions with AWS Lambda functions to process the files. The developer used Amazon API Gateway and Lambda functions to create an API to support the UI.
The company's Ul team reports that the request to process a file is often returning timeout errors because of the see or complexity of the files. The Ul team wants the API to provide an immediate response so that the Ul can deploy a message while the files are being processed. The backend process that is invoked by the API needs to send an email message when the report processing is complete.
What should the developer do to configure the API to meet these requirements?

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

Answer: A

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

NEW QUESTION 9
A company is building a serverless application on AWS. The application uses an AWS Lambda function to process customer orders 24 hours a day, 7 days a week. The Lambda function calls an external vendor's HTTP API to process payments.
During load tests, a developer discovers that the external vendor payment processing API occasionally times out and returns errors. The company expects that some payment processing API calls will return errors.
The company wants the support team to receive notifications in near real time only when the payment processing external API error rate exceed 5% of the total number of transactions in an hour. Developers need to use an existing Amazon Simple Notification Service (Amazon SNS) topic that is configured to notify the support team.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 10
A company is building an application for stock trading. The application needs sub- millisecond latency for processing trade requests. The company uses Amazon DynamoDB to store all the trading data that is used to process each trading request A development team performs load testing on the application and finds that the data retrieval time is higher than expected. The development team needs a solution that reduces the data retrieval time with the least possible effort.
Which solution meets these requirements'?

A. Add local secondary indexes (LSis) for the trading data.
B. Store the trading data m Amazon S3 and use S3 Transfer Acceleration.
C. Add retries with exponential back off for DynamoDB queries.
D. Use DynamoDB Accelerator (DAX) to cache the trading data.

Answer: D

Explanation:
This solution will meet the requirements by using DynamoDB Accelerator (DAX), which is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10 times performance improvement - from milliseconds to microseconds - even at millions of requests per second. The developer can use DAX to cache the trading data that is used to process each trading request, which will reduce the data retrieval time with the least possible effort. Option A is not optimal because it will add local secondary indexes (LSIs) for the trading data, which may not improve the performance or reduce the latency of data retrieval, as LSIs are stored on the same partition as the base table and share the same provisioned throughput. Option B is not optimal because it will store the trading data in Amazon S3 and use S3 Transfer Acceleration, which is a feature that enables fast, easy, and secure transfers of files over long distances between S3 buckets and clients, not between DynamoDB and clients. Option C is not optimal because it will add retries with exponential backoff for DynamoDB queries, which is a strategy to handle transient errors by retrying failed requests with increasing delays, not by reducing data retrieval time.
References: [DynamoDB Accelerator (DAX)], [Local Secondary Indexes]

NEW QUESTION 10
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 15
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 20
A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node is application. To minimize these bugs, the developer wants to impendent automated testing of Lambda functions in an environment that Closely simulates the Lambda environment.
The developer needs to give other developers the ability to run the tests locally. The developer also needs to integrate the tests into the team's continuous integration and continuous delivery (Ct/CO) pipeline before the AWS Cloud Development Kit (AWS COK) deployment.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 21
A developer is configuring an applications deployment environment in AWS CodePipeine. The application code is stored in a GitHub repository. The developer wants to ensure that the repository package's unit tests run in the new deployment environment. The deployment has already set the pipeline's source provider to GitHub and has specified the repository and branch to use in the deployment.
When combination of steps should the developer take next to meet these requirements with the least the LEAST overhead' (Select TWO).

A. Create an AWS CodeCommt projec
B. Add the repository package's build and test commands to the protects buildspec
C. Create an AWS CodeBuid projec
commands to the projects buildspec
D. Add the repository package's build and test
E. Create an AWS CodeDeploy protec
F. Add the repository package's build and test commands to the project's buildspec
G. Add an action to the source stag
H. Specify the newly created project as the action provide
I. Specify the build attract as the actions input artifact.
J. Add a new stage to the pipeline alter the source stag
K. Add an action to the new stag
L. Speedy the newly created protect as the action provide
M. Specify the source artifact as the action's input artifact.

Answer: BE

Explanation:
This solution will ensure that the repository package’s unit tests run in the new deployment environment with the least overhead because it uses AWS CodeBuild to build and test the code in a fully managed service, and AWS CodePipeline to orchestrate the deployment stages and actions. Option A is not optimal because it will use AWS CodeCommit instead of AWS CodeBuild, which is a source control service, not a build and test service. Option C is not optimal because it will use AWS CodeDeploy instead of AWS CodeBuild, which is a deployment service, not a build and test service. Option D is not optimal because it will add an action to the source stage instead of creating a new stage, which will not follow the best practice of separating different deployment phases. References: AWS CodeBuild, AWS CodePipeline

NEW QUESTION 22
A company has an application that is hosted on Amazon EC2 instances The application stores objects in an Amazon S3 bucket and allows users to download

objects from the S3 bucket A developer turns on S3 Block Public Access for the S3 bucket After this change, users report errors when they attempt to download objects. The developer needs to implement a solution so that only users who are signed in to the application can access objects in the S3 bucket.
Which combination of steps will meet these requirements in the MOST secure way? (Select TWO.)

A. Create an EC2 instance profile and role with an appropriate policy Associate the role with the EC2 instances
B. Create an 1AM user with an appropriate polic
C. Store the access key ID and secret access key on the EC2 instances
D. Modify the application to use the S3 GeneratePresignedUrl API call
E. Modify the application to use the S3 GetObject API call and to return the object handle to the user
F. Modify the application to delegate requests to the S3 bucket.

Answer: AC

Explanation:
The most secure way to allow the EC2 instances to access the S3 bucket is to use an EC2 instance profile and role with an appropriate policy that grants the necessary permissions. This way, the EC2 instances can use temporary security credentials that are automatically rotated and do not need to store any access keys on the instances. To allow the users who are signed in to the application to download objects from the S3 bucket, the application can use the S3 GeneratePresignedUrl API call to create a pre-signed URL that grants temporary access to a specific object. The pre-signed URL can be returned to the user, who can then use it to download the object within a specified time period. References
? Use Amazon S3 with Amazon EC2
? How to Access AWS S3 Bucket from EC2 Instance In a Secured Way
? Sharing an Object with Others

NEW QUESTION 23
A company uses Amazon API Gateway to expose a set of APIs to customers. The APIs have caching enabled in API Gateway. Customers need a way to invalidate the cache for each API when they test the API.
What should a developer do to give customers the ability to invalidate the API cache?

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

Answer: D

NEW QUESTION 26
A developer is creating an AWS Lambda function that searches for Items from an Amazon DynamoDQ table that contains customer contact information. The DynamoDB table items have the customers as the partition and additional properties such as customer -type, name, and job_title.
The Lambda function runs whenever a user types a new character into the customer_type text Input. The developer wants to search to return partial matches of all tne 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 input, as the partition key and email_address as the sort ke
B. Perform a query operation on the GSI by using the begins with key condition expression with the email_address property.
C. Add a global secondary index (GSI) to the DynamoDB table with email_address as the partition key and customer_type as the sort ke
D. Perform a query operation on the GSI by using the begine_with key condition expresses with the emai
E. Address property.
F. Add a local secondary index (LSI) to the DynemoOB table with customer_type as the partition Key and email_address as the sort Ke
G. Perform a quick operation on the LSI by using the begine_with Key condition expression with the email-address property.
H. Add a local secondary index (LSI) to the DynamoDB table with job-title as the partition key and email_address as the sort ke
I. Perform a query operation on the LSI by using the begins_with key condition expression with the email_address property.

Answer: A

Explanation:
The solution that will meet the requirements is to 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 begins_with key condition expression with the email_address property. This way, the developer can search for partial matches of the email_address property of a particular customer type without recreating the DynamoDB table. The other options either involve using a local secondary index (LSI), which requires recreating the table, or using a different partition key, which does not allow filtering by customer_type.
Reference: Using Global Secondary Indexes in DynamoDB

NEW QUESTION 31
A developer wants to expand an application to run in multiple AWS Regions. The developer wants to copy Amazon Machine Images (AMIs) with the latest changes and create a new application stack in the destination Region. According to company requirements, all AMIs must be encrypted in all Regions. However, not all the AMIs that the company uses are encrypted.
How can the developer expand the application to run in the destination Region while meeting the encryption requirement?

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Amazon Machine Images (AMIs) are encrypted snapshots of EC2 instances that can be used to launch new instances. The developer can create new AMIs from the existing instances and specify encryption parameters. The developer can copy the encrypted AMIs to the destination Region and use them to create a new

application stack. The developer can delete the unencrypted AMIs after the encryption process is complete. This solution will meet the encryption requirement and allow the developer to expand the application to run in the destination Region.
References:
? [Amazon Machine Images (AMI) - Amazon Elastic Compute Cloud]
? [Encrypting an Amazon EBS Snapshot - Amazon Elastic Compute Cloud]
? [Copying an AMI - Amazon Elastic Compute Cloud]

NEW QUESTION 32
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 37
A developer is creating an AWS Lambda function that consumes messages from an Amazon Simple Queue Service (Amazon SQS) standard queue. The developer notices that the Lambda function processes some messages multiple times.
How should developer resolve this issue MOST cost-effectively?

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

Answer: A

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

NEW QUESTION 39
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 42
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 47
A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

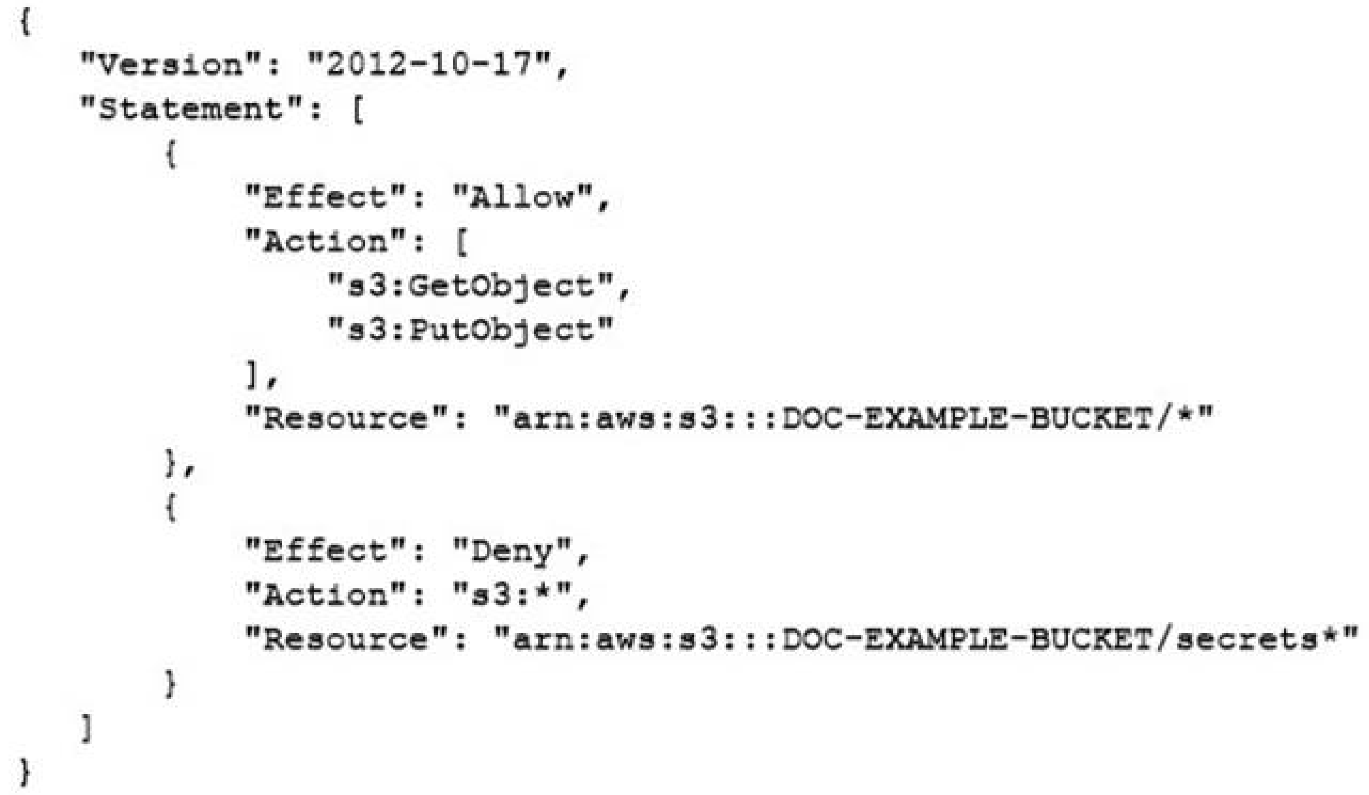

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetObject",
                "s3:PutObject"
            ],
            "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
        },
        {
            "Effect": "Deny",
            "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 51
A company receives food orders from multiple partners. The company has a microservices application that uses Amazon API Gateway APIs with AWS Lambda integration. Each partner sends orders by calling a customized API that is exposed through API Gateway. The API call invokes a shared Lambda function to process the orders.
Partners need to be notified after the Lambda function processes the orders. Each partner must receive updates for only the partner's own orders. The company wants to add new partners in the future with the fewest code changes possible.
Which solution will meet these requirements in the MOST scalable way?

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

Answer: C

Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fully managed messaging service that enables pub/sub communication between distributed systems. The developer can create an SNS topic and configure the Lambda function to publish messages with specific attributes to the topic. The developer can subscribe each partner to the SNS topic and apply the appropriate filter policy to the topic subscriptions. This way, each partner will receive updates for only their own orders based on the message attributes. This solution will meet the requirements in the most scalable way and allow adding new partners in the future with minimal code changes.
References:
? [Amazon Simple Notification Service (SNS)]
? [Filtering Messages with Attributes - Amazon Simple Notification Service]

NEW QUESTION 55
A developer is creating an application that will be deployed on IoT devices. The application will send data to a RESTful API that is deployed as an AWS Lambda function. The application will assign each API request a unique identifier. The volume of API requests from the application can randomly increase at any given time of day.
During periods of request throttling, the application might need to retry requests. The API must be able to handle duplicate requests without inconsistencies or data loss.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 60
A company has an existing application that has hardcoded database credentials A developer needs to modify the existing application The application is deployed in two AWS Regions with an active-passive failover configuration to meet company's disaster recovery strategy
The developer needs a solution to store the credentials outside the code. The solution must comply With the company's disaster recovery strategy
Which solution Will meet these requirements in the MOST secure way?

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

Answer: A

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

NEW QUESTION 64
A company is building a web application on AWS. When a customer sends a request, the application will generate reports and then make the reports available to the customer within one hour. Reports should be accessible to the customer for 8 hours. Some reports are larger than 1 MB. Each report is unique to the customer. The application should delete all reports that are older than 2 days.
Which solution will meet these requirements with the LEAST operational overhead?

A. Generate the reports and then store the reports as Amazon DynamoDB items that have a specified TT
B. Generate a URL that retrieves the reports from DynamoD

C. Provide the URL to customers through the web application.
D. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryptio
E. Attach the reports to an Amazon Simple Notification Service (Amazon SNS) messag
F. Subscribe the customer to email notifications from Amazon SNS.
G. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryptio
H. Generate a presigned URL that contains an expiration date Provide the URL to customers through the web applicatio
I. Add S3 Lifecycle configuration rules to the S3 bucket to delete old reports.
J. Generate the reports and then store the reports in an Amazon RDS database with a date stam
K. Generate an URL that retrieves the reports from the RDS databas
L. Provide the URL to customers through the web applicatio
M. Schedule an hourly AWS Lambda function to delete database records that have expired date stamps.

Answer: C

Explanation:
This solution will meet the requirements with the least operational overhead because it uses Amazon S3 as a scalable, secure, and durable storage service for the reports. The presigned URL will allow customers to access their reports for a limited time (8 hours) without requiring additional authentication. The S3 Lifecycle configuration rules will automatically delete the reports that are older than 2 days, reducing storage costs and complying with the data retention policy. Option A is not optimal because it will incur additional costs and complexity to store the reports as DynamoDB items, which have a size limit of 400 KB. Option B is not optimal because it will not provide customers with access to their reports within one hour, as Amazon SNS email delivery is not guaranteed. Option D is not optimal because it will require more operational overhead to manage an RDS database and a Lambda function for storing and deleting the reports.
References: Amazon S3 Presigned URLs, Amazon S3 Lifecycle

NEW QUESTION 67
For a deployment using AWS Code Deploy, what is the run order of the hooks for in-place deployments?

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

Answer: B

Explanation:
For in-place deployments, AWS CodeDeploy uses a set of predefined hooks that run in a specific order during each deployment lifecycle event. The hooks are ApplicationStop, BeforeInstall, AfterInstall, ApplicationStart, and ValidateService. The run order of the hooks for in-place deployments is as follows:
? ApplicationStop: This hook runs first on all instances and stops the current application that is running on the instances.
? BeforeInstall: This hook runs after ApplicationStop on all instances and performs any tasks required before installing the new application revision.
? AfterInstall: This hook runs after BeforeInstall on all instances and performs any tasks required after installing the new application revision.
? ApplicationStart: This hook runs after AfterInstall on all instances and starts the new application that has been installed on the instances.
? ValidateService: This hook runs last on all instances and verifies that the new application is running properly on the instances.
Reference: [AWS CodeDeploy lifecycle event hooks reference]

NEW QUESTION 69
A company is migrating an on-premises database to Amazon RDS for MySQL. The company has read-heavy workloads. The company wants to refactor the code to achieve optimum read performance for queries.
Which solution will meet this requirement with LEAST current and future effort?

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

Answer: C

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

NEW QUESTION 74
A company has an application that stores data in Amazon RDS instances. The application periodically experiences surges of high traffic that cause performance problems.
During periods of peak traffic, a developer notices a reduction in query speed in all database queries.
The team's technical lead determines that a multi-threaded and scalable caching solution should be used to offload the heavy read traffic. The solution needs to improve performance.
Which solution will meet these requirements with the LEAST complexity?

A. Use Amazon ElastiCache for Memcached to offload read requests from the main database.
B. Replicate the data to Amazon DynamoD

C. Set up a DynamoDB Accelerator (DAX) cluster.
D. Configure the Amazon RDS instances to use Multi-AZ deployment with one standby instanc
E. Offload read requests from the main database to the standby instance.
F. Use Amazon ElastiCache for Redis to offload read requests from the main database.

Answer: A

Explanation:
? Amazon ElastiCache for Memcached is a fully managed, multithreaded, and scalable in-memory key-value store that can be used to cache frequently accessed data and improve application performance1. By using Amazon ElastiCache for Memcached, the developer can reduce the load on the main database and handle high traffic surges more efficiently.
? To use Amazon ElastiCache for Memcached, the developer needs to create a cache cluster with one or more nodes, and configure the application to store and retrieve data from the cache cluster2. The developer can use any of the supported Memcached clients to interact with the cache cluster3. The developer can also use Auto Discovery to dynamically discover and connect to all cache nodes in a cluster4.
? Amazon ElastiCache for Memcached is compatible with the Memcached protocol, which means that the developer can use existing tools and libraries that work with Memcached1. Amazon ElastiCache for Memcached also supports data partitioning, which allows the developer to distribute data among multiple nodes and scale out the cache cluster as needed.
? Using Amazon ElastiCache for Memcached is a simple and effective solution that meets the requirements with the least complexity. The developer does not need to change the database schema, migrate data to a different service, or use a different caching model. The developer can leverage the existing Memcached ecosystem and easily integrate it with the application.

NEW QUESTION 76
A developer is trying get data from an Amazon DynamoDB table called demoman-table. The developer configured the AWS CLI to use a specific IAM use's credentials and ran the following command.

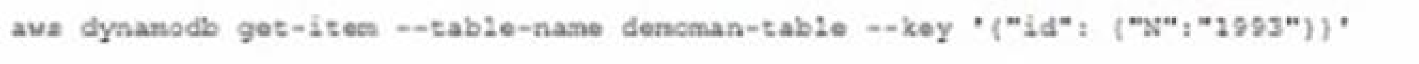

```
aws dynamodb get-item --table-name demoman-table --key '{"id": {"N":"1993"}}'
```

The command returned errors and no rows were returned. What is the MOST likely cause of these issues?

A. The command is incorrect; it should be rewritten to use put-item with a string argument
B. The developer needs to log a ticket with AWS Support to enable access to the demoman-table
C. Amazon DynamoOB cannot be accessed from the AWS CLI and needs to called via the REST API
D. The IAM user needs an associated policy with read access to demoman-table

Answer: D

Explanation:
This solution will most likely solve the issues because it will grant the IAM user the necessary permission to access the DynamoDB table using the AWS CLI command. The error message indicates that the IAM user does not have sufficient access rights to perform the scan operation on the table. Option A is not optimal because it will change the command to use put-item instead of scan, which will not achieve the desired result of getting data from the table. Option B is not optimal because it will involve contacting AWS Support, which may not be necessary or efficient for this issue. Option C is not optimal because it will state that DynamoDB cannot be accessed from the AWS CLI, which is incorrect as DynamoDB supports AWS CLI commands.
References: AWS CLI for DynamoDB, [IAM Policies for DynamoDB]

NEW QUESTION 80
An application that runs on AWS Lambda requires access to specific highly confidential objects in an Amazon S3 bucket. In accordance with the principle of least privilege a company grants access to the S3 bucket by using only temporary credentials.
How can a developer configure access to the S3 bucket in the MOST secure way?

A. Hardcode the credentials that are required to access the S3 objects in the application cod
B. Use the credentials to access me required S3 objects.
C. Create a secret access key and access key ID with permission to access the S3 bucke
D. Store the key and key ID in AWS Secrets Manage
E. Configure the application to retrieve the Secrets Manager secret and use the credentials to access me S3 objects.
F. Create a Lambda function execution role Attach a policy to the rote that grants access to specific objects in the S3 bucket.
G. Create a secret access key and access key ID with permission to access the S3 bucket Store the key and key ID as environment variables m Lambd
H. Use the environment variables to access the required S3 objects.

Answer: C

Explanation:
This solution will meet the requirements by creating a Lambda function execution role, which is an IAM role that grants permissions to a Lambda function to access AWS resources such as Amazon S3 objects. The developer can attach a policy to the role that grants access to specific objects in the S3 bucket that are required by the application, following the principle of least privilege. Option A is not optimal because it will hardcode the credentials that are required to access S3 objects in the application code, which is insecure and difficult to maintain. Option B is not optimal because it will create a secret access key and access key ID with permission to access the S3 bucket, which will introduce additional security risks and complexity for storing and managing credentials. Option D is not optimal because it will store the secret access key and access key ID as environment variables in Lambda, which is also insecure and difficult to maintain. References: [AWS Lambda Execution Role], [Using AWS Lambda with Amazon S3]

NEW QUESTION 84
A company is running Amazon EC2 instances in multiple AWS accounts. A developer needs to implement an application that collects all the lifecycle events of the EC2 instances. The application needs to store the lifecycle events in a single Amazon Simple Queue Service (Amazon SQS) queue in the company's main AWS account for further processing.
Which solution will meet these requirements?

A. Configure Amazon EC2 to deliver the EC2 instance lifecycle events from all accounts to the Amazon EventBridge event bus of the main accoun
B. Add an EventBridge rule to the event bus of the main account that matches all EC2 instance lifecycle event
C. Add the SQS queue as a target of the rule.
D. Use the resource policies of the SQS queue in the main account to give each account permissions to write to that SQS queu

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

Answer: D

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

NEW QUESTION 86
A developer is deploying an AWS Lambda function The developer wants the ability to return to older versions of the function quickly and seamlessly.
How can the developer achieve this goal with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 91
A developer has created an AWS Lambda function that makes queries to an Amazon Aurora MySQL DB instance. When the developer performs a test the OB instance shows an error for too many connections.
Which solution will meet these requirements with the LEAST operational effort?

A. Create a read replica for the DB instance Query the replica DB instance instead of the primary DB instance.
B. Migrate the data lo an Amazon DynamoDB database.
C. Configure the Amazon Aurora MySQL DB instance tor Multi-AZ deployment.
D. Create a proxy in Amazon RDS Proxy Query the proxy instead of the DB instance.

Answer: D

Explanation:
This solution will meet the requirements by using Amazon RDS Proxy, which is a fully managed, highly available database proxy for Amazon RDS that makes applications more scalable, more resilient to database failures, and more secure. The developer can create a proxy in Amazon RDS Proxy, which sits between the application
and the DB instance and handles connection management, pooling, and routing. The developer can query the proxy instead of the DB instance, which reduces the number of open connections to the DB instance and avoids errors for too many connections. Option A is not optimal because it will create a read replica for the DB instance, which may not solve the problem of too many connections as read replicas also have connection limits and may incur additional costs. Option B is not optimal because it will migrate the data to an Amazon DynamoDB database, which may introduce additional complexity and overhead for migrating and accessing data from a different database service. Option C is not optimal because it will configure the Amazon Aurora MySQL DB instance for Multi-AZ deployment, which may improve availability and durability of the DB instance but not reduce the number of connections.
References: [Amazon RDS Proxy], [Working with Amazon RDS Proxy]

NEW QUESTION 96
A company has an analytics application that uses an AWS Lambda function to process transaction data asynchronously A developer notices that asynchronous invocations of the Lambda function sometimes fail When failed Lambda function invocations occur, the developer wants to invoke a second Lambda function to handle errors and log details.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

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

invocations. References
? Using AWS Lambda destinations
? Asynchronous invocation - AWS Lambda
? AWS Lambda Destinations: What They Are and Why to Use Them
? AWS Lambda Destinations: A Complete Guide | Dashbird

NEW QUESTION 98
A developer is using AWS Amplify Hosting to build and deploy an application. The developer is receiving an increased number of bug reports from users. The developer wants to add end-to-end testing to the application to eliminate as many bugs as possible before the bugs reach production.
Which solution should the developer implement to meet these requirements?

A. Run the amplify add test command in the Amplify CLI.
B. Create unit tests in the applicatio
C. Deploy the unit tests by using the amplify push command in the Amplify CLI.
D. Add a test phase to the amplify.yml build settings for the application.
E. Add a test phase to the aws-exports.js file for the application.

Answer: C

Explanation:
The solution that will meet the requirements is to add a test phase to the amplify.yml build settings for the application. This way, the developer can run end-to-end tests on every code commit and catch any bugs before deploying to production. The other options either do not support end-to-end testing, or do not run tests automatically.
Reference: End-to-end testing

NEW QUESTION 101
A developer uses AWS CloudFormation to deploy an Amazon API Gateway API and an AWS Step Functions state machine The state machine must reference the API Gateway API after the CloudFormation template is deployed The developer needs a solution that uses the state machine to reference the API Gateway endpoint.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: A

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

NEW QUESTION 103
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 106
A developer deployed an application to an Amazon EC2 instance The application needs to know the public IPv4 address of the instance
How can the application find this information?

A. Query the instance metadata from http./M69.254.169.254. latestmeta-data/.
B: Query the instance user data from http '169 254.169 254. latest/user-data/
C. Query the Amazon Machine Image (AMI) information from http://169.254.169.254/latest/meta-data/ami/.
D. Check the hosts file of the operating system

Answer: A

Explanation:
The instance metadata service provides information about the EC2 instance, including the public IPv4 address, which can be obtained by querying the endpoint http://169.254.169.254/latest/meta-data/public-ipv4. References
? Instance metadata and user data
? Get Public IP Address on current EC2 Instance
? Get the public ip address of your EC2 instance quickly

NEW QUESTION 111
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 115
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 120
A developer warns to add request validation to a production environment Amazon API Gateway API. The developer needs to test the changes before the API is deployed to the production environment. For the lest the developer will send test requests to the API through a testing tool.
Which solution will meet these requirements with the LEAST operational overhead?

A. Export the existing API to an OpenAPI fil

B. Create a new API Import the OpenAPI file Modify the new API to add request validatio
C. Perform the tests Modify the existing API to add request validatio
D. Deploy the existing API to production.
E. Modify the existing API to add request validatio
F. Deploy the updated API to a new API Gateway stage Perform the tests Deploy the updated API to the API Gateway production stage.
G. Create a new API Add the necessary resources and methods including new request validatio
H. Perform the tests Modify the existing API to add request validatio
I. Deploy the existing API to production.
J. Clone the exiting API Modify the new API lo add request validatio
Modify the existing API to add request validation Deploy the existing API to production.
K. Perform the tests

Answer: D

Explanation:
This solution allows the developer to test the changes without affecting the production environment. Cloning an API creates a copy of the API definition that can be modified independently. The developer can then add request validation to the new API and test it using a testing tool. After verifying that the changes work as expected, the developer can apply the same changes to the existing API and deploy it to production.
Reference: Clone an API, [Enable Request Validation for an API in API Gateway]

NEW QUESTION 121
A developer migrated a legacy application to an AWS Lambda function. The function uses a third-party service to pull data with a series of API calls at the end of each month. The function than processes the data to generate the monthly reports. The function has Been working with no issues so far.
The third-party service recently issued a restriction to allow a feed number to API calls each minute and each day. If the API calls exceed the limit tor each minute or each day, then the service will produce errors. The API also provides the minute limit and daily limit in the response header. This restriction might extend the overall process to multiple days because the process is consuming more API calls than the available limit.
What is the MOST operationally efficient way to refactor the server less application to accommodate this change?

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

Answer: A

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

NEW QUESTION 122
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 125
A company is using Amazon RDS as the Backend database for its application. After a recent marketing campaign, a surge of read requests to the database increased the latency of data retrieval from the database.
The company has decided to implement a caching layer in front of the database. The cached content must be encrypted and must be highly available.
Which solution will meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 130

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 135

A developer is building a web application that uses Amazon API Gateway to expose an AWS Lambda function to process requests from clients. During testing, the developer notices that the API Gateway times out even though the Lambda function finishes under
the set time limit.
Which of the following API Gateway metrics in Amazon CloudWatch can help the developer troubleshoot the issue? (Choose two.)

A. CacheHitCount
B. IntegrationLatency
C. CacheMissCount
D. Latency
E. Count

Answer: BD

Explanation:

Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. Amazon CloudWatch is a service that monitors AWS resources and applications. API Gateway provides several CloudWatch metrics to help developers troubleshoot issues with their APIs. Two of the metrics that can help the developer troubleshoot the issue of API Gateway timing out are:
? IntegrationLatency: This metric measures the time between when API Gateway
relays a request to the backend and when it receives a response from the backend. A high value for this metric indicates that the backend is taking too long to respond and may cause API Gateway to time out.
? Latency: This metric measures the time between when API Gateway receives a
request from a client and when it returns a response to the client. A high value for this metric indicates that either the integration latency is high or API Gateway is taking too long to process the request or response.
References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [Amazon API Gateway Metrics and Dimensions - Amazon CloudWatch]
? [Troubleshooting API Errors - Amazon API Gateway]

NEW QUESTION 139

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

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

Answer: D

Explanation:

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

NEW QUESTION 143

A developer is designing a serverless application with two AWS Lambda functions to process photos. One Lambda function stores objects in an Amazon S3 bucket and stores the associated metadata in an Amazon DynamoDB table. The other Lambda function fetches the objects from the S3 bucket by using the metadata from the DynamoDB table. Both Lambda functions use the same Python library to perform complex computations and are approaching the quota for the

maximum size of zipped deployment packages.
What should the developer do to reduce the size of the Lambda deployment packages with the LEAST operational overhead?

A. Package each Python library in its own .zip file archiv
B. Deploy each Lambda function with its own copy of the library.
C. Create a Lambda layer with the required Python librar
D. Use the Lambda layer in both Lambda functions.
E. Combine the two Lambda functions into one Lambda functio
F. Deploy the Lambda function as a single .zip file archive.
G. Download the Python library to an S3 bucke
H. Program the Lambda functions to reference the object URLs.

Answer: B

Explanation:
AWS Lambda is a service that lets developers run code without provisioning or managing servers. Lambda layers are a distribution mechanism for libraries, custom runtimes, and other dependencies. The developer can create a Lambda layer with the required Python library and use the layer in both Lambda functions. This will reduce the size of the Lambda deployment packages and avoid reaching the quota for the maximum size of zipped deployment packages. The developer can also benefit from using layers to manage dependencies separately from function code.
References:
? [What Is AWS Lambda? - AWS Lambda]
? [AWS Lambda Layers - AWS Lambda]

NEW QUESTION 145

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

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

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 148
A company is planning to use AWS CodeDeploy to deploy an application to Amazon Elastic Container Service (Amazon ECS) During the deployment of a new version of the application, the company initially must expose only 10% of live traffic to the new version of the deployed application. Then, after 15 minutes elapse, the company must route all the remaining live traffic to the new version of the deployed application.
Which CodeDeploy predefined configuration will meet these requirements?

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

Answer: A

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

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

NEW QUESTION 152
A developer is creating an AWS Lambda function. The Lambda function needs an external library to connect to a third-party solution The external library is a collection of files with a total size of 100 MB The developer needs to make the external library available to the Lambda execution environment and reduce the Lambda package space
Which solution will meet these requirements with the LEAST operational overhead?

A.

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

Answer: A

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

NEW QUESTION 157
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 162
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 166
A company developed an API application on AWS by using Amazon CloudFront. Amazon API Gateway, and AWS Lambda. The API has a minimum of four requests every second A developer notices that many API users run the same query by using the POST method. The developer wants to cache the POST request to optimize the API resources.
Which solution will meet these requirements'?

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

Answer: A

Explanation:
This solution will meet the requirements by using Amazon CloudFront, which is a content delivery network (CDN) service that speeds up the delivery of web content and APIs to end users. The developer can configure the CloudFront cache, which is a set of edge locations that store copies of popular or recently accessed content close to the viewers. The developer can also update the application to return cached content based upon the default request headers, which are a set of HTTP headers that CloudFront automatically forwards to the origin server and uses to determine whether an object in an edge location is still valid. By caching the POST requests, the developer can optimize the API resources and reduce the latency for repeated queries. Option B is not optimal because it will override the cache method in the selected stage of API Gateway, which is not possible or effective as API Gateway does not support caching for POST methods by default. Option C is not optimal because it will save the latest request response in Lambda /tmp directory, which is a local storage space that is available for each Lambda function invocation, not a cache that can be shared across multiple invocations or requests. Option D is not optimal because it will save the latest request in AWS Systems Manager Parameter Store, which is a service that provides secure and scalable storage for configuration data and secrets, not a cache for API responses.
References: [Amazon CloudFront], [Caching Content Based on Request Headers]

NEW QUESTION 169

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

A. Migrate the secret credentials to Amazon RDS parameter group
B. Encrypt the parameter by using an AWS Key Management Service (AWS KMS) key Turn on secret rotatio
C. Use 1AM policies and roles to grant AWS KMS permissions to access Amazon RDS.
D. Migrate the credentials to AWS Systems Manager Parameter Stor
E. Encrypt the parameter by using an AWS Key Management Service (AWS KMS) ke
F. Turn on secret rotatio
G. Use 1AM policies and roles to grant Amazon ECS Fargate permissions to access to AWS Secrets Manager
H. Migrate the credentials to ECS Fargate environment variable
I. Encrypt the credentials by using an AWS Key Management Service (AWS KMS) key Turn on secret rotatio
J. Use 1AM policies and roles to grant Amazon ECS Fargate permissions to access to AWS Secrets Manager.
K. Migrate the credentials to AWS Secrets Manage
L. Encrypt the credentials by using an AWS Key Management Service (AWS KMS) key Turn on secret rotation Use 1AM policies and roles to grant Amazon ECS Fargate permissions to access to AWS Secrets Manager by using keys.

Answer: D

Explanation:
AWS Secrets Manager is a service that helps you store, distribute, and rotate secrets securely. You can use Secrets Manager to migrate your credentials from your application code to a secure and encrypted storage. You can also enable automatic rotation of your secrets by using AWS Lambda functions or custom logic. You can use IAM policies and roles to grant your Amazon ECS Fargate tasks permissions to access your secrets from Secrets Manager. This solution minimizes the operational overhead of managing your credentials and enhances the security of your application. References
? AWS Secrets Manager: Store, Distribute, and Rotate Credentials Securely | AWS
News Blog
? Why You Should Audit and Rotate Your AWS Credentials Periodically - Cloud Academy
? Top 5 AWS root account best practices - TheServerSide

NEW QUESTION 171
A company has deployed an application on AWS Elastic Beanstalk. The company has configured the Auto Scaling group that is associated with the Elastic Beanstalk environment to have five Amazon EC2 instances. If the capacity is fewer than four EC2 instances during the deployment, application performance degrades. The company is using the all-at-once deployment policy.
What is the MOST cost-effective way to solve the deployment issue?

A. Change the Auto Scaling group to six desired instances.
B. Change the deployment policy to traffic splittin
C. Specify an evaluation time of 1 hour.
D. Change the deployment policy to rolling with additional batc
E. Specify a batch size of 1.
F. Change the deployment policy to rollin
G. Specify a batch size of 2.

Answer: C

Explanation:
This solution will solve the deployment issue by deploying the new version of the application to one new EC2 instance at a time, while keeping the old version running on

the existing instances. This way, there will always be at least four instances serving traffic during the deployment, and no downtime or performance degradation will occur. Option A is not optimal because it will increase the cost of running the Elastic Beanstalk environment without solving the deployment issue. Option B is not optimal because it will split the traffic between two versions of the application, which may cause inconsistency and confusion for the customers. Option D is not optimal because it will deploy the new version of the application to two existing instances at a time, which may reduce the capacity below four instances during the deployment.
References: AWS Elastic Beanstalk Deployment Policies

NEW QUESTION 173

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

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

Answer: A

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

NEW QUESTION 178
An application that is deployed to Amazon EC2 is using Amazon DynamoDB. The app cation calls the DynamoDB REST API Periodically the application receives a ProvisionedThroughputExceededException error when the application writes to a DynamoDB table.
Which solutions will mitigate this error MOST cost-effectively^ (Select TWO)

A. Modify the application code to perform exponential back off when the error is received.
B. Modify the application to use the AWS SDKs for DynamoDB.

C. Increase the read and write throughput of the DynamoDB table.
D. Create a DynamoDB Accelerator (DAX) cluster for the DynamoDB table.
E. Create a second DynamoDB table Distribute the reads and writes between the two tables.

Answer: AB

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

NEW QUESTION 183
A company has an Amazon S3 bucket containing premier content that it intends to make available to only paid subscribers of its website. The S3 bucket currently has default permissions of all objects being private to prevent inadvertent exposure of the premier content to non-paying website visitors.
How can the company Limit the ability to download a premier content file in the S3 Bucket to paid subscribers only?

A. Apply a bucket policy that allows anonymous users to download the content from the S3 bucket.
B. Generate a pre-signed object URL for the premier content file when a pad subscriber requests a download.
C. Add a Docket policy that requires multi-factor authentication for request to access the S3 bucket objects.
D. Enable server-side encryption on the S3 bucket for data protection against the non- paying website visitors.

Answer: B

Explanation:
This solution will limit the ability to download a premier content file in the S3 bucket to paid subscribers only because it uses a pre-signed object URL that grants temporary access to an S3 object for a specified duration. The pre-signed object URL can be generated by the company's website when a paid subscriber requests a download, and can be verified by Amazon S3 using the signature in the URL. Option A is not optimal because it will allow anyone to download the content from the S3 bucket without verifying their subscription status. Option C is not optimal because it will require additional steps and costs to configure multi-factor authentication for accessing the S3 bucket objects, which may not be feasible or user-friendly for paid subscribers. Option D is not optimal because it will not prevent non-paying website visitors from accessing the S3 bucket objects, but only encrypt them at rest.
References: Share an Object with Others, [Using Amazon S3 Pre-Signed URLs]

NEW QUESTION 185
A company is building a scalable data management solution by using AWS services to improve the speed and agility of development. The solution will ingest large volumes of data from various sources and will process this data through multiple business rules and transformations.
The solution requires business rules to run in sequence and to handle reprocessing of data if errors occur when the business rules run. The company needs the solution to be scalable and to require the least possible maintenance.
Which AWS service should the company use to manage and automate the orchestration of the data flows to meet these requirements?

A. AWS Batch
B. AWS Step Functions
C.

AWS Glue
D. AWS Lambda

Answer: B

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

NEW QUESTION 189
A developer is designing an AWS Lambda function that creates temporary files that are less than 10 MB during invocation. The temporary files will be accessed and modified multiple times during invocation. The developer has no need to save or retrieve these files in the future.
Where should the temporary files be stored?

A. the /tmp directory
B. Amazon Elastic File System (Amazon EFS)
C. Amazon Elastic Block Store (Amazon EBS)
D. Amazon S3

Answer: A

Explanation:
AWS Lambda is a service that lets developers run code without provisioning or managing servers. Lambda provides a local file system that can be used to store temporary files during invocation. The local file system is mounted under the /tmp directory and has a limit of 512 MB. The temporary files are accessible only by the Lambda function that created them and are deleted after the function execution ends. The developer can store temporary files that are less than 10 MB in the /tmp directory and access and modify them multiple times during invocation.
References:
? [What Is AWS Lambda? - AWS Lambda]
? [AWS Lambda Execution Environment - AWS Lambda]

NEW QUESTION 194
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 195
A company wants to automate part of its deployment process. A developer needs to automate the process of checking for and deleting unused resources that supported previously deployed stacks but that are no longer used.

The company has a central application that uses the AWS Cloud Development Kit (AWS CDK) to manage all deployment stacks. The stacks are spread out across multiple accounts. The developer's solution must integrate as seamlessly as possible within the current deployment process.
Which solution will meet these requirements with the LEAST amount of configuration?

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

Answer: B

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

NEW QUESTION 196
......