Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

https://www.2passeasy.com/dumps/AWS-Certified-Developer-Associate/

NEW QUESTION 1
An Amazon Simple Queue Service (Amazon SQS) queue serves as an event source for an AWS Lambda function In the SQS queue, each item corresponds to a video file that the Lambda function must convert to a smaller resolution The Lambda function is timing out on longer video files, but the Lambda function's timeout is already configured to its maximum value
What should a developer do to avoid the timeouts without additional code changes'?

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

Answer: A

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

NEW QUESTION 2
A company notices that credentials that the company uses to connect to an external software as a service (SaaS) vendor are stored in a configuration file as plaintext.
The developer needs to secure the API credentials and enforce automatic credentials rotation on a quarterly basis.
Which solution will meet these requirements MOST securely?

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

Answer: C

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

NEW QUESTION 3
A company needs to deploy all its cloud resources by using AWS CloudFormation templates A developer must create an Amazon Simple Notification Service (Amazon SNS) automatic notification to help enforce this rule. The developer creates an SNS topic and subscribes the email address of the company's security team to the SNS topic.
The security team must receive a notification immediately if an 1AM role is created without the use of CloudFormation.
Which solution will meet this requirement?

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

Answer: D

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

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

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

Answer: B

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

NEW QUESTION 5
A developer is creating an AWS Lambda function that needs credentials to connect to an Amazon RDS for MySQL database. An Amazon S3 bucket currently stores the credentials. The developer needs to improve the existing solution by implementing credential rotation and secure storage. The developer also needs to provide integration with the Lambda function.
Which solution should the developer use to store and retrieve the credentials with the LEAST management overhead?

A. Store the credentials in AWS Systems Manager Parameter Stor
B. Select the database that the parameter will acces
C. Use the default AWS Key Management Service (AWS KMS) key to encrypt the paramete
D. Enable automatic rotation for the paramete
E. Use the parameter from Parameter Store on the Lambda function to connect to the database.
F. Encrypt the credentials with the default AWS Key Management Service (AWS KMS) ke
G. Store the credentials as environment variables for the Lambda functio
H. Create a second Lambda function to generate new credentials and to rotate the credentials by updating the environment variables of the first Lambda functio
I. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedul
J. Update the database to use the new credential
K. On the first Lambda function, retrieve the credentials from the environment variable
L. Decrypt the credentials by using AWS KMS, Connect to the database.
M. Store the credentials in AWS Secrets Manage
N. Set the secret type to Credentials for Amazon RDS databas
O. Select the database that the secret will acces
P. Use the default AWS Key Management Service (AWS KMS) key to encrypt the secre
Q. Enable automatic rotation for the secre
R. Use the secret from Secrets Manager on the Lambda function to connect to the database.
S. Encrypt the credentials by using AWS Key Management Service (AWS KMS). Store the credentials in an Amazon DynamoDB tabl
T. Create a second Lambda function to rotate the credential
. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedul
. Update the DynamoDB tabl
. Update the database to use the generated credential
. Retrieve the credentials from DynamoDB with the first Lambda functio
. Connect to the database.

Answer: C

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

NEW QUESTION 6
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 7
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 8
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 9
A company needs to harden its container images before the images are in a running state. The company's application uses Amazon Elastic Container Registry (Amazon ECR) as an image registry. Amazon Elastic Kubernetes Service (Amazon EKS) for compute, and an AWS CodePipeline pipeline that orchestrates a continuous integration and continuous delivery (CI/CD) workflow.
Dynamic application security testing occurs in the final stage of the pipeline after a new image is deployed to a development namespace in the EKS cluster. A developer needs to
place an analysis stage before this deployment to analyze the container image earlier in the CI/CD pipeline.
Which solution will meet these requirements with the MOST operational efficiency?

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

Answer: B

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

NEW QUESTION 10
A developer designed an application on an Amazon EC2 instance The application makes API requests to objects in an Amazon S3 bucket
Which combination of steps will ensure that the application makes the API requests in the MOST secure manner? (Select TWO.)

A. Create an IAM user that has permissions to the S3 bucke
B. Add the user to an 1AM group
C. Create an IAM role that has permissions to the S3 bucket
D. Add the IAM role to an instance profil
E. Attach the instance profile to the EC2 instance.
F. Create an 1AM role that has permissions to the S3 bucket Assign the role to an 1AM group
G. Store the credentials of the IAM user in the environment variables on the EC2 instance

Answer: BC

Explanation:
- Create an IAM role that has permissions to the S3 bucket. - Add the IAM role to an instance profile. Attach the instance profile to the EC2 instance. We first need to create a n IAM Role with permissions to read and eventually write a specific S3 bucket. Then, we need to attach the role to the EC2 isntance through an instance profile. In this
way, the ec2 instance has the permissions to read and eventually write the specified S3 bucket

NEW QUESTION 10
A developer is working on a Python application that runs on Amazon EC2 instances. The developer wants to enable tracing of application requests to debug performance issues in the code.
Which combination of actions should the developer take to achieve this goal? (Select TWO)

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

Answer: BE

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

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

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

Answer: BC

Explanation:
The UnprocessedKeys element indicates that the BatchGetItem operation did not process all of the requested items in the current response. This can happen if the
response size limit is exceeded or if the table's provisioned throughput is exceeded. To handle this situation, the developer should retry the batch operation with exponential backoff and randomized delay to avoid throttling errors and reduce the load on the table. The developer should also use an AWS SDK to make the requests, as the SDKs automatically retry requests that return UnprocessedKeys.
References:
? [BatchGetItem - Amazon DynamoDB]
? [Working with Queries and Scans - Amazon DynamoDB]
? [Best Practices for Handling DynamoDB Throttling Errors]

NEW QUESTION 15
A developer is working on an ecommerce platform that communicates with several third- party payment processing APIs The third-party payment services do not provide a test environment.
The developer needs to validate the ecommerce platform's integration with the third-party payment processing APIs. The developer must test the API integration code without invoking the third-party payment processing APIs.
Which solution will meet these requirements'?

A. Set up an Amazon API Gateway REST API with a gateway response configured for status code 200 Add response templates that contain sample responses captured from the real third-party API.

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

Answer: D

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

NEW QUESTION 20
An developer is building a serverless application by using the AWS Serverless Application Model (AWS SAM). The developer is currently testing the application in a development environment. When the application is nearly finsihed, the developer will need to set up additional testing and staging environments for a quality assurance team.
The developer wants to use a feature of the AWS SAM to set up deployments to multiple environments.
Which solution will meet these requirements with the LEAST development effort?

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

Answer: A

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

NEW QUESTION 21
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?

A. Use a multi-AZ Amazon RDS deploymen
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 26
A developer at a company needs to create a small application that makes the same API call once each day at a designated time. The company does not have infrastructure in the AWS Cloud yet, but the company wants to implement this functionality on AWS.
Which solution meets these requirements in the MOST operationally efficient manner?

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

Answer: C

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

NEW QUESTION 27
An application uses Lambda functions to extract metadata from files uploaded to an S3 bucket; the metadata is stored in Amazon DynamoDB. The application starts behaving unexpectedly, and the developer wants to examine the logs of the Lambda function code for errors.
Based on this system configuration, where would the developer find the logs?

A. Amazon S3
B. AWS CloudTrail
C. Amazon CloudWatch
D. Amazon DynamoDB

Answer: C

Explanation:
Amazon CloudWatch is the service that collects and stores logs from AWS Lambda functions. The developer can use CloudWatch Logs Insights to query and analyze the logs for errors and metrics. Option A is not correct because Amazon S3 is a storage service that does not store Lambda function logs. Option B is not correct because AWS CloudTrail is a service that records API calls and events for AWS services, not Lambda function logs. Option D is not correct because Amazon DynamoDB is a database service that does not store Lambda function logs.
References: AWS Lambda Monitoring, [CloudWatch Logs Insights]

NEW QUESTION 32
A developer is migrating some features from a legacy monolithic application to use AWS Lambda functions instead. The application currently stores data in an Amazon Aurora DB cluster that runs in private subnets in a VPC. The AWS account has one VPC deployed. The Lambda functions and the DB cluster are deployed in the same AWS Region in the same AWS account.
The developer needs to ensure that the Lambda functions can securely access the DB cluster without crossing the public internet.
Which solution will meet these requirements?

A. Configure the DB cluster's public access setting to Yes.

B. Configure an Amazon RDS database proxy for the Lambda functions.
C. Configure a NAT gateway and a security group for the Lambda functions.
D. Configure the VPC, subnets, and a security group for the Lambda functions.

Answer: D

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

NEW QUESTION 34
A company has multiple Amazon VPC endpoints in the same VPC. A developer needs configure an Amazon S3 bucket policy so users can access an S3 bucket only by using these VPC endpoints.
Which solution will meet these requirements?

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

Answer: D

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

NEW QUESTION 39
A company is implementing an application on Amazon EC2 instances. The application needs to process incoming transactions. When the application detects a transaction that is not valid, the application must send a chat message to the company's support team. To send the message, the application needs to retrieve the access token to authenticate by using the chat API.
A developer needs to implement a solution to store the access token. The access token must be encrypted at rest and in transit. The access token must also be accessible from other AWS accounts.
Which solution will meet these requirements with the LEAST management overhead?

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

Answer: C

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

NEW QUESTION 41
A developer is creating an AWS Lambda function that searches for items from an Amazon DynamoDB table that contains customer contact information- The DynamoDB table items have the customer's email_address as the partition key and additional properties such as customer_type, name, and job_tltle.

The Lambda function runs whenever a user types a new character into the customer_type text input The developer wants the search to return partial matches of all the email_address property of a particular customer_type The developer does not want to recreate the DynamoDB table.
What should the developer do to meet these requirements?

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

Answer: A

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

NEW QUESTION 43
A company has installed smart motes in all Its customer locations. The smart meter's measure power usage at 1minute intervals and send the usage readings to a remote endpoint tot collection. The company needs to create an endpoint that will receive the smart meter readings and store the readings in a database. The company wants to store the location ID and timestamp information.
The company wants to give Is customers low-latency access to their current usage and historical usage on demand The company expects demand to increase significantly. The solution must not impact performance or include downtime write seeing.
When solution will meet these requirements MOST cost-effectively?

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

Answer: B

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

NEW QUESTION 46
A company is creating an application that processes csv files from Amazon S3 A developer has created an S3 bucket The developer has also created an AWS Lambda function to process the csv files from the S3 bucket
Which combination of steps will invoke the Lambda function when a csv file is uploaded to Amazon S3? (Select TWO.)

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

Answer: AC

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

NEW QUESTION 51
A developer needs to store configuration variables for an application. The developer needs to set an expiration date and time for me configuration. The developer wants to receive notifications. Before the configuration expires. Which solution will meet these requirements with the LEAST operational overhead?

A. Create a standard parameter in AWS Systems Manager Parameter Store Set Expiation and Expiration Notification policy types.
B. Create a standard parameter in AWS Systems Manager Parameter Store Create an AWS Lambda function to expire the configuration and to send Amazon Simple Notification Service (Amazon SNS) notifications.
C. Create an advanced parameter in AWS Systems Manager Parameter Store Set Expiration and Expiration Notification policy types.
D. Create an advanced parameter in AWS Systems Manager Parameter Store Create an Amazon EC2 instance with a corn job to expire the configuration and to send notifications.

Answer: C

Explanation:
This solution will meet the requirements by creating an advanced parameter in AWS Systems Manager Parameter Store, which is a secure and scalable service for storing and managing configuration data and secrets. The advanced parameter allows setting expiration and expiration notification policy types, which enable specifying an expiration date and time for the configuration and receiving notifications before the configuration expires. 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 create a standard parameter in AWS Systems Manager Parameter Store, which does not support expiration and expiration notification policy types. 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 create a Docker container from Node.js base image to invoke Lambda functions, which will incur additional costs and overhead for creating and running Docker containers. References: AWS Systems Manager Parameter Store, [Using SSL/TLS to Encrypt a Connection to a DB Instance]

NEW QUESTION 54
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

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

Answer: B

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

NEW QUESTION 56
A developer is troubleshooting an application mat uses Amazon DynamoDB in the uswest- 2 Region. The application is deployed to an Amazon EC2 instance. The application requires read-only permissions to a table that is named Cars The EC2 instance has an attached IAM role that contains the following IAM policy.

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

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

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

Answer: C

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

NEW QUESTION 59
A developer is preparing to begin development of a new version of an application. The previous version of the application is deployed in a production environment. The developer needs to deploy fixes and updates to the current version during the development of the new version of the application. The code for the new version of the application is stored in AWS CodeCommit.
Which solution will meet these requirements?

A. From the main branch, create a feature branch for production bug fixe
B. Create a second feature branch from the main branch for development of the new version.
C. Create a Git tag of the code that is currently deployed in productio
D. Create a Git tag for the development of the new versio
E. Push the two tags to the CodeCommit repository.

F. From the main branch, create a branch of the code that is currently deployed in productio
G. Apply an IAM policy that ensures no other other users can push or merge to the branch.
H. Create a new CodeCommit repository for development of the new version of the applicatio
I. Create a Git tag for the development of the new version.

Answer: A

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

NEW QUESTION 61
A developer is creating a service that uses an Amazon S3 bucket for image uploads. The service will use an AWS Lambda function to create a thumbnail of each image Each time an image is uploaded the service needs to send an email notification and create the thumbnail The developer needs to configure the image processing and email notifications setup.
Which solution will meet these requirements?

A. Create an Amazon Simple Notification Service (Amazon SNS) topic Configure S3 event notifications with a destination of the SNS topic Subscribe the Lambda function to the SNS topic Create an email notification subscription to the SNS topic
B. Create an Amazon Simple Notification Service (Amazon SNS) topi
C. Configure S3 event notifications with a destination of the SNS topi
D. Subscribe the Lambda function to the SNS topi
E. Create an Amazon Simple Queue Service (Amazon SQS) queue Subscribe the SQS queue to the SNS topic Create an email notification subscription to the SQS queue.
F. Create an Amazon Simple Queue Service (Amazon SQS) queue Configure S3 event notifications with a destination of the SQS queue Subscribe the Lambda function to the SQS queue Create an email notification subscription to the SQS queue.
G. Create an Amazon Simple Queue Service (Amazon SQS) queu
H. Send S3 event notifications to Amazon EventBridg
I. Create an EventBndge rule that runs the Lambda function when images are uploaded to the S3 bucket Create an EventBridge rule that sends notifications to the SQS queue Create an email notification subscription to the SQS queue

Answer: A

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

NEW QUESTION 66
A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function.
How should the developer configure the Lambda function to detect changes to the DynamoDB table?

A. Create an Amazon Kinesis data stream, and attach it to the DynamoDB tabl
B. Create a trigger to connect the data stream to the Lambda function.
C. Create an Amazon EventBridge rule to invoke the Lambda function on a regular schedul
D. Connect to the DynamoDB table from the Lambda function to detect changes.
E. Enable DynamoDB Streams on the tabl
F. Create a trigger to connect the DynamoDB stream to the Lambda function.
G. Create an Amazon Kinesis Data Firehose delivery stream, and attach it to the DynamoDB tabl
H. Configure the delivery stream destination as the Lambda function.

Answer: C

Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance with seamless scalability. DynamoDB Streams is a feature that captures data modification events in DynamoDB tables. The developer can enable DynamoDB Streams on the table and create a trigger to connect the DynamoDB stream to the Lambda function. This solution will enable the Lambda function to detect changes to the DynamoDB table in near real time.
References:
? [Amazon DynamoDB]
? [DynamoDB Streams - Amazon DynamoDB]
? [Using AWS Lambda with Amazon DynamoDB - AWS Lambda]

NEW QUESTION 67
A company is running Amazon EC2 instances in multiple AWS accounts. A developer needs to implement an application that collects all the lifecycle events of the EC2 instances. The application needs to store the lifecycle events in a single Amazon Simple Queue Service (Amazon SQS) queue in the company's main AWS account for further processing.
Which solution will meet these requirements?

A. Configure Amazon EC2 to deliver the EC2 instance lifecycle events from all accounts to the Amazon EventBridge event bus of the main accoun

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

Answer: D

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

NEW QUESTION 72
A company has an ecommerce application. To track product reviews, the company's development team uses an Amazon DynamoDB table.
Every record includes the following
• A Review ID a 16-digrt universally unique identifier (UUID)
• A Product ID and User ID 16 digit UUIDs that reference other tables
• A Product Rating on a scale of 1-5
• An optional comment from the user
The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product.
Which index will provide the FASTEST response for this query"?

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

Answer: A

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

NEW QUESTION 75
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 78
A developer is working on an ecommerce website The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available
How can the developer update the application to meet these requirements with MINIMUM changes?

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

Answer: D

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

NEW QUESTION 83
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 cach
B. Update the application to return cached content based upon the default request headers.
C. Override the cache method in the selected stage of API Gatewa
D. Select the POST method.
E. Save the latest request response in Lambda /tmp director
F. Update the Lambda function to check the /tmp directory.
G. Save the latest request in AWS Systems Manager Parameter Stor
H. Modify the Lambda function to take the latest request response from Parameter Store.

Answer: B

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

NEW QUESTION 85

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 88
A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the AWS Key Management Service (AWS KMS) key was used and by whom.
Which encryption option will meet these requirements?

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

Answer: B

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

NEW QUESTION 93
A developer wants to add request validation to a production environment Amazon API Gateway API. The developer needs to test the changes before the API is deployed to the production environment. For the test, the developer will send test requests to the API through a testing tool.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 97

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

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

Answer: A

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

NEW QUESTION 100
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 105
A company is offering APIs as a service over the internet to provide unauthenticated read access to statistical information that is updated daily. The company uses Amazon API Gateway and AWS Lambda to develop the APIs. The service has become popular, and the company wants to enhance the responsiveness of the APIs.
Which action can help the company achieve this goal?

A. Enable API caching in API Gateway.
B. Configure API Gateway to use an interface VPC endpoint.
C. Enable cross-origin resource sharing (CORS) for the APIs.
D. Configure usage plans and API keys in API Gateway.

Answer: A

Explanation:
Amazon API Gateway is a service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. The developer can enable API caching in API Gateway to cache responses from the backend integration point for a specified time-to- live (TTL) period. This can improve the responsiveness of the APIs by reducing the number of calls made to the backend service. References:
? [What Is Amazon API Gateway? - Amazon API Gateway]
? [Enable API Caching to Enhance Responsiveness - Amazon API Gateway]

NEW QUESTION 106
A developer is optimizing an AWS Lambda function and wants to test the changes in

production on a small percentage of all traffic. The Lambda function serves requests to a REST API in Amazon API Gateway. The developer needs to deploy their changes and perform a test in production without changing the API Gateway URL.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 107
A developer is designing a serverless application for a game in which users register and log in through a web browser The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API
The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.
Which solution will meet these requirements'?

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

Answer: A

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

NEW QUESTION 110
An Amazon Kinesis Data Firehose delivery stream is receiving customer data that contains personally identifiable information. A developer needs to remove pattern-based customer identifiers from the data and store the modified data in an Amazon S3 bucket.
What should the developer do to meet these requirements?

A. Implement Kinesis Data Firehose data transformation as an AWS Lambda functio
B. Configure the function to remove the customer identifier
C. Set an Amazon S3 bucket as the destination of the delivery stream.
D. Launch an Amazon EC2 instanc
E. Set the EC2 instance as the destination of the delivery strea
F. Run an application on the EC2 instance to remove the customer identifier
G. Store the transformed data in an Amazon S3 bucket.
H. Create an Amazon OpenSearch Service instanc
I. Set the OpenSearch Service instance as the destination of the delivery strea
J. Use search and replace to remove the customer identifier
K. Export the data to an Amazon S3 bucket.
L. Create an AWS Step Functions workflow to remove the customer identifier
M. As the last step in the workflow, store the transformed data in an Amazon S3 bucke
N. Set the workflow as the destination of the delivery stream.

Answer: A

Explanation:
Amazon Kinesis Data Firehose is a service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon OpenSearch

Service, and Amazon Kinesis Data Analytics. The developer can implement Kinesis Data Firehose data transformation as an AWS Lambda function. The function can remove pattern-based customer identifiers from the data and return the modified data to Kinesis Data Firehose. The developer can set an Amazon S3 bucket as the destination of the delivery stream. References:
? [What Is Amazon Kinesis Data Firehose? - Amazon Kinesis Data Firehose]
? [Data Transformation - Amazon Kinesis Data Firehose]

NEW QUESTION 115
An application uses an Amazon EC2 Auto Scaling group. A developer notices that EC2 instances are taking a long time to become available during scale-out events. The UserData script is taking a long time to run.
The developer must implement a solution to decrease the time that elapses before an EC2 instance becomes available. The solution must make the most recent version of the application available at all times and must apply all available security updates. The solution also must minimize the number of images that are created. The images must be validated.
Which combination of steps should the developer take to meet these requirements? (Choose two.)

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

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

NEW QUESTION 120

A developer wants to store information about movies. Each movie has a title, release year, and genre. The movie information also can include additional properties about the cast and production crew. This additional information is inconsistent across movies. For example, one movie might have an assistant director, and another movie might have an animal trainer.
The developer needs to implement a solution to support the following use cases:
For a given title and release year, get all details about the movie that has that title and release year.
For a given title, get all details about all movies that have that title. For a given genre, get all details about all movies in that genre. Which data store configuration will meet these requirements?

A. Create an Amazon DynamoDB tabl
B. Configure the table with a primary key that consists of the title as the partition key and the release year as the sort ke
C. Create a global secondary index that uses the genre as the partition key and the title as the sort key.
D. Create an Amazon DynamoDB tabl
E. Configure the table with a primary key that consists of the genre as the partition key and the release year as the sort ke
F. Create a global secondary index that uses the title as the partition key.
G. On an Amazon RDS DB instance, create a table that contains columns for title, release year, and genr
H. Configure the title as the primary key.
I. On an Amazon RDS DB instance, create a table where the primary key is the title and all other data is encoded into JSON format as one additional column.

Answer: A

Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and consistent performance with seamless scalability. The developer can create a DynamoDB table and configure the table with a primary key that consists of the title as the partition key and the release year as the sort key. This will enable querying for a given title and release year efficiently. The developer can also create a global secondary index that uses the genre as the partition key and the title as the sort key. This will enable querying for a given genre efficiently. The developer can store additional properties about the cast and production crew as attributes in the DynamoDB table. These attributes can have different data types and structures, and they do not need to be consistent across items.
References:
? [Amazon DynamoDB]
? [Working with Queries - Amazon DynamoDB]
? [Working with Global Secondary Indexes - Amazon DynamoDB]

NEW QUESTION 125

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 129

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 130
A company is using Amazon API Gateway to invoke a new AWS Lambda function The company has Lambda function versions in its PROD and DEV environments. In each environment, there is a Lambda function alias pointing to the corresponding Lambda function version API Gateway has one stage that is configured to point at the PROD alias
The company wants to configure API Gateway to enable the PROD and DEV Lambda function versions to be simultaneously and distinctly available
Which solution will meet these requirements?

A. Enable a Lambda authorizer for the Lambda function alias in API Gateway Republish PROD and create a new stage for DEV Create API Gateway stage variables for the PROD and DEV stage
B. Point each stage variable to the PROD Lambda authorizer to the DEV Lambda authorizer.
C. Set up a gateway response in API Gateway for the Lambda function alia
D. Republish PROD and create a new stage for DE
E. Create gateway responses in API Gateway for PROD and DEV Lambda aliases
F. Use an environment variable for the Lambda function alias in API Gatewa
G. Republish PROD and create a new stage for developmen
H. Create API gateway environment variables for PROD and DEV stage
I. Point each stage variable to the PROD Lambda function alias to the DEV Lambda function alias.
J. Use an API Gateway stage variable to configure the Lambda function alias Republish PROD and create a new stage for development Create API Gateway stage variables for PROD and DEV stages Point each stage variable to the PROD Lambda function alias and to the DEV Lambda function alias

Answer: D

Explanation:
The best solution is to use an API Gateway stage variable to configure the Lambda function alias. This allows you to specify the Lambda function name and its alias or version using the syntax function_name:$ {stageVariables.variable_name} in the Integration Request. You can then create different stages in API Gateway, such as PROD and DEV, and assign different values to the stage variable for each stage. This way, you can invoke different Lambda function versions or aliases based on the stage that you are using, without changing the function name in the Integration Request. References
? Using API Gateway stage variables to manage Lambda functions
? How to point AWS API gateway stage to specific lambda function alias?
? Setting stage variables using the Amazon API Gateway console
? Amazon API Gateway stage variables reference

NEW QUESTION 135
When using the AWS Encryption SDK how does the developer keep track of the data encryption keys used to encrypt data?

A. The developer must manually keep Hack of the data encryption keys used for each data object.
B. The SDK encrypts the data encryption key and stores it (encrypted) as part of the resumed ophertext.
C. The SDK stores the data encryption keys automaticity in Amazon S3.
D. The data encryption key is stored m the user data for the EC2 instance.

Answer: B

Explanation:
This solution will meet the requirements by using AWS Encryption SDK, which is a client-side encryption library that enables developers to encrypt and decrypt data using data encryption keys that are protected by AWS Key Management Service (AWS KMS). The SDK encrypts the data encryption key with a customer master key (CMK) that is managed by AWS KMS, and stores it (encrypted) as part of the returned ciphertext. The developer does not need to keep track of the data encryption keys used to encrypt data, as they are stored with the encrypted data and can be retrieved and decrypted by using AWS KMS when needed.
Option A is not optimal because it will require manual tracking of the data encryption keys used for each data object, which is error-prone and inefficient. Option C is not optimal because it will store the data encryption keys automatically in Amazon S3, which is unnecessary and insecure as Amazon S3 is not designed for storing encryption keys. Option D is not optimal because it will store the data encryption key in the user data for the EC2 instance, which is also unnecessary and insecure as user data is not encrypted by default.
References: [AWS Encryption SDK], [AWS Key Management Service]

NEW QUESTION 138

A developer is modifying an existing AWS Lambda function White checking the code the developer notices hardcoded parameter various for an Amazon RDS for SQL Server user name password database host and port. There also are hardcoded parameter values for an Amazon DynamoOB table. an Amazon S3 bucket, and an Amazon Simple Notification Service (Amazon SNS) topic.
The developer wants to securely store the parameter values outside the code m an encrypted format and wants to turn on rotation for the credentials. The developer also wants to be able to reuse the parameter values from other applications and to update the parameter values without modifying code.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an RDS database secret in AWS Secrets Manage
B. Set the user name password, database, host and por
C. Turn on secret rotatio
D. Create encrypted Lambda environment variables for the DynamoDB table, S3 bucket and SNS topic.
E. Create an RDS database secret in AWS Secrets Manage
F. Set the user name password, database, host and por
G. Turn on secret rotatio
H. Create Secure String parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket and SNS topic.
I. Create RDS database parameters in AWS Systems Manager Paramete
J. Store for the user name password, database, host and por
K. Create encrypted Lambda environment variables for me DynamoDB table, S3 bucket, and SNS topi
L. Create a Lambda function and set the logic for the credentials rotation task Schedule the credentials rotation task in Amazon EventBridge.
M. Create RDS database parameters in AWS Systems Manager Paramete
N. Store for the user name password database, host, and por
O. Store the DynamoDB tabl
P. S3 bucket, and SNS topic in Amazon S3 Create a Lambda function and set the logic for the credentials rotation Invoke the Lambda function on a schedule.

Answer: B

Explanation:
This solution will meet the requirements by using AWS Secrets Manager and AWS Systems Manager Parameter Store to securely store the parameter values outside the code in an encrypted format. AWS Secrets Manager is a service that helps protect secrets such as database credentials by encrypting them with AWS Key Management Service (AWS KMS) and enabling automatic rotation of secrets. The developer can create an RDS database secret in AWS Secrets Manager

and set the user name, password, database, host, and port for accessing the RDS database. The developer can also turn on secret rotation, which will change the database credentials periodically according to a specified schedule or event. AWS Systems Manager Parameter Store is a service that provides secure and scalable storage for configuration data and secrets. The developer can create Secure String parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket, and SNS topic, which will encrypt them with AWS KMS. The developer can also reuse the parameter values from other applications and update them without modifying code. Option A is not optimal because it will create encrypted Lambda environment variables for the DynamoDB table, S3 bucket, and SNS topic, which may not be reusable or updatable without modifying code. Option C is not optimal because it will create RDS database parameters in AWS Systems Manager Parameter Store, which does not support automatic rotation of secrets. Option D is not optimal because it will store the DynamoDB table, S3 bucket, and SNS topic in Amazon S3, which may introduce additional costs and complexity for accessing configuration data.
References: AWS Secrets Manager, [AWS Systems Manager Parameter Store]

NEW QUESTION 143
A company has developed a new serverless application using AWS Lambda functions that will be deployed using the AWS Serverless Application Model (AWS SAM) CLI.
Which step should the developer complete prior to deploying the application?

A. Compress the application to a zip file and upload it into AWS Lambda.
B. Test the new AWS Lambda function by first tracing it m AWS X-Ray.
C. Bundle the serverless application using a SAM package.
D. Create the application environment using the eb create my-env command.

Answer: C

Explanation:
This step should be completed prior to deploying the application because it prepares the application artifacts for deployment. The AWS Serverless Application Model (AWS SAM) is a framework that simplifies building and deploying serverless applications on AWS. The AWS SAM CLI is a command-line tool that helps you create, test, and deploy serverless applications using AWS SAM templates. The sam package command bundles the application artifacts, such as Lambda function code and API definitions, and uploads them to an Amazon S3 bucket. The command also returns a CloudFormation template that is ready to be deployed with the sam deploy command. Compressing the application to a zip file and uploading it to AWS Lambda will not work because it does not use AWS SAM templates or CloudFormation. Testing the new Lambda function by first tracing it in AWS X- Ray will not prepare the application for deployment, but only monitor its performance and errors. Creating the application environment using the eb create my-env command will not work because it is a command for AWS Elastic Beanstalk, not AWS SAM.

NEW QUESTION 147
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 148
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Developer-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Developer-Associate Product From:

https://www.2passeasy.com/dumps/AWS-Certified-Developer-Associate/

Money Back Guarantee

AWS-Certified-Developer-Associate Practice Exam Features:

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

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

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

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