Exam Questions DVA-C02

DVA-C02

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

NEW QUESTION 1
A data visualization company wants to strengthen the security of its core applications The applications are deployed on AWS across its development staging, pre-production, and production environments. The company needs to encrypt all of its stored sensitive credentials The sensitive credentials need to be automatically rotated Aversion of the sensitive credentials need to be stored for each environment
Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: D

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

NEW QUESTION 2
A developer is creating a mobile app that calls a backend service by using an Amazon API Gateway REST API. For integration testing during the development phase, the developer wants to simulate different backend responses without invoking the backend service.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AWS Lambda functio
B. Use API Gateway proxy integration to return constant HTTP responses.
C. Create an Amazon EC2 instance that serves the backend REST API by using an AWS CloudFormation template.
D. Customize the API Gateway stage to select a response type based on the request.
E. Use a request mapping template to select the mock integration response.

Answer: D

Explanation:
Amazon API Gateway supports mock integration responses, which are predefined responses that can be returned without sending requests to a backend service. Mock integration responses can be used for testing or prototyping purposes, or for simulating different backend responses based on certain conditions. A request mapping template can be used to select a mock integration response based on an expression that evaluates some aspects of the request, such as headers, query strings, or body content. This solution does not require any additional resources or code changes and has the least operational overhead. Reference: Set up mock integrations for an API Gateway REST API
https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-mock- integration.html

NEW QUESTION 3
A developer is deploying a new application to Amazon Elastic Container Service (Amazon ECS). The developer needs to securely store and retrieve different types of variables. These variables include authentication information for a remote API, the URL for the API, and credentials. The authentication information and API URL must be available to all current and future deployed versions of the application across development, testing, and production environments.
How should the developer retrieve the variables with the FEWEST application changes?

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

Answer: A

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

NEW QUESTION 4
A developer maintains a critical business application that uses Amazon DynamoDB as the primary data store The DynamoDB table contains millions of documents and receives 30- 60 requests each minute The developer needs to perform processing in near-real time on the documents when they are added or updated in the DynamoDB table

How can the developer implement this feature with the LEAST amount of change to the existing application code?

A. Set up a cron job on an Amazon EC2 instance Run a script every hour to query the table for changes and process the documents
B. Enable a DynamoDB stream on the table Invoke an AWS Lambda function to process the documents.
C. Update the application to send a PutEvents request to Amazon EventBridg
D. Create an EventBridge rule to invoke an AWS Lambda function to process the documents.
E. Update the application to synchronously process the documents directly after the DynamoDB write

Answer: B

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

NEW QUESTION 5
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 6
A company has an application that runs across multiple AWS Regions. The application is experiencing performance issues at irregular intervals. A developer must
                                   to implement distributed tracing for the application to troubleshoot the root cause of the performance issues.
use AWS X-Ray
What should the developer do to meet this requirement?

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

Answer: B

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

NEW QUESTION 7
A developer needs to migrate an online retail application to AWS to handle an anticipated increase in traffic. The application currently runs on two servers: one server for the web application and another server for the database. The web server renders webpages and manages session state in memory. The database server hosts a MySQL database that contains order details. When traffic to the application is heavy, the memory usage for the web server approaches 100% and the application slows down considerably.
The developer has found that most of the memory increase and performance decrease is related to the load of managing additional user sessions. For the web server migration, the developer will use Amazon EC2 instances with an Auto Scaling group behind an Application Load Balancer.
Which additional set of changes should the developer make to the application to improve the application's performance?

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

Answer: B

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

NEW QUESTION 8

A developer is testing a new file storage application that uses an Amazon CloudFront distribution to serve content from an Amazon S3 bucket. The distribution accesses the S3 bucket by using an origin access identity (OAI). The S3 bucket's permissions explicitly deny access to all other users. The application prompts users to authenticate on a login page and then uses signed cookies to allow users to access their personal storage directories. The developer has configured the distribution to use its default cache behavior with restricted viewer access and has set the origin to point to the S3 bucket. However, when the developer tries to navigate to the login page, the developer receives a 403 Forbidden error.

The developer needs to implement a solution to allow unauthenticated access to the login page. The solution also must keep all private content secure.
Which solution will meet these requirements?

A. Add a second cache behavior to the distribution with the same origin as the default cache behavio
B. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricte
C. Keep the default cache behavior's settings unchanged.
D. Add a second cache behavior to the distribution with the same origin as the default cache behavio
E. Set the path pattern for the second cache behavior to *, and make viewer access restricte
F. Change the default cache behavior's path pattern to the path of the login page, and make viewer access unrestricted.
G. Add a second origin as a failover origin to the default cache behavio
H. Point the failover origin to the S3 bucke
I. Set the path pattern for the primary origin to *, and make viewer access restricte
J. Set the path pattern for the failover origin to the path of the login page, and make viewer access unrestricted.
K. Add a bucket policy to the S3 bucket to allow read acces
L. Set the resource on the policy to the Amazon Resource Name (ARN) of the login page object in the S3 bucke
M. Add a CloudFront function to the default cache behavior to redirect unauthorized requests to the login page's S3 URL.

Answer: A

Explanation:
The solution that will meet the requirements is to add a second cache behavior to the distribution with the same origin as the default cache behavior. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricted. Keep the default cache behavior's settings unchanged. This way, the login page can be accessed without authentication, while all other content remains secure and requires signed cookies. The other options either do not allow unauthenticated access to the login page, or expose private content to unauthorized users.
Reference: Restricting Access to Amazon S3 Content by Using an Origin Access Identity

NEW QUESTION 9

An online food company provides an Amazon API Gateway HTTP API 1o receive orders for partners. The API is integrated with an AWS Lambda function. The Lambda function stores the orders in an Amazon DynamoDB table.

The company expects to onboard additional partners Some to me panthers require additional Lambda function to receive orders. The company has created an Amazon S3 bucket. The company needs 10 store all orders and updates m the S3 bucket for future analysis
How can the developer ensure that an orders and updates are stored to Amazon S3 with the LEAST development effort?

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

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

Answer: C

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

NEW QUESTION 10

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

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

Answer:

D

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

NEW QUESTION 10
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 13
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 14
A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node is application. To minimize these bugs, the developer wants to impendent automated testing of Lambda functions in an environment that Closely simulates the Lambda environment.
The developer needs to give other developers the ability to run the tests locally. The developer also needs to integrate the tests into the team's continuous integration and continuous delivery (Ct/CO) pipeline before the AWS Cloud Development Kit (AWS COK) deployment.
Which solution will meet these requirements?

A. Create sample events based on the Lambda documentatio
B. Create automated test scripts that use the cdk local invoke command to invoke the Lambda function
C. Check the response Document the test scripts for the other developers on the team Update the CI/CD pipeline to run the test scripts.

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

Answer: C

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

NEW QUESTION 18
A developer is configuring an applications deployment environment in AWS CodePipeine. The application code is stored in a GitHub repository. The developer wants to ensure that the repository package's unit tests run in the new deployment environment. The deployment has already set the pipeline's source provider to GitHub and has specified the repository and branch to use in the deployment.
When combination of steps should the developer take next to meet these requirements with the least the LEAST overhead' (Select TWO).

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

Answer: BE

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

NEW QUESTION 22
A company has an application that is hosted on Amazon EC2 instances The application stores objects in an Amazon S3 bucket and allows users to download objects from the S3 bucket A developer turns on S3 Block Public Access for the S3 bucket After this change, users report errors when they attempt to download objects The developer needs to implement a solution so that only users who are signed in to the application can access objects in the S3 bucket.
Which combination of steps will meet these requirements in the MOST secure way? (Select TWO.)

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

Answer: AC

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

NEW QUESTION 27
A company has a web application that is hosted on Amazon EC2 instances The EC2 instances are configured to stream logs to Amazon CloudWatch Logs The company needs to receive an Amazon Simple Notification Service (Amazon SNS) notification when the number of application error messages exceeds a defined threshold within a 5-minute period
Which solution will meet these requirements?

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

Answer: D

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

NEW QUESTION 31
An online sales company is developing a serverless application that runs on AWS. The application uses an AWS Lambda function that calculates order success rates and stores the data in an Amazon DynamoDB table. A developer wants an efficient way to invoke the Lambda function every 15 minutes.
Which solution will meet this requirement with the LEAST development effort?

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

Answer: A

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

NEW QUESTION 34
A developer is creating an application that will store personal health information (PHI). The PHI needs to be encrypted at all times. An encrypted Amazon RDS for MySQL DB instance is storing the data. The developer wants to increase the performance of the application by caching frequently accessed data while adding the ability to sort or rank the cached datasets.
Which solution will meet these requirements?

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

Answer: A

Explanation:
 Amazon ElastiCache is a service that offers fully managed in-memory data stores that are compatible with Redis or Memcached. The developer can create an ElastiCache for Redis instance and enable encryption of data in transit and at rest. This will ensure that the PHI is encrypted at all times. The developer can store frequently accessed data in the cache and use Redis features such as sorting and ranking to enhance the performance of the application.
References:
? [What Is Amazon ElastiCache? - Amazon ElastiCache]
                          ? [Encryption in Transit - Amazon ElastiCache for Redis]
? [Encryption at Rest - Amazon ElastiCache for Redis]

NEW QUESTION 38

A developer is creating a new REST API by using Amazon API Gateway and AWS Lambda. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.

The developer wants to make the REST API available for testing by using API Gateway locally.

Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

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

Answer: D

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

NEW QUESTION 39
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 44
A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream. The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is increasing and that the Lambda run duration is constantly above normal.
Which actions should the developer take to increase the processing speed? (Choose two.)

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

Answer: AC

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

NEW QUESTION 48
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 49
A company uses a custom root certificate authority certificate chain (Root CA Cert) that is 10 KB in size generate SSL certificates for its on-premises HTTPS
endpoints. One of the company's cloud based applications has hundreds of AWS Lambda functions that pull date from these endpoints. A developer updated the
trust store of the Lambda execution environment to use the Root CA Cert when the Lambda execution environment is initialized. The developer bundled the Root
CA Cert as a text file in the Lambdas deployment bundle.
After 3 months of development the root CA Cert is no longer valid and must be updated. The developer needs a more efficient solution to update the Root CA Cert
for all deployed Lambda functions. The solution must not include rebuilding or updating all Lambda functions that use the Root CA Cert. The solution must also
work for all development, testing and production environment. Each environment is managed in a separate AWS account.
When combination of steps Would the developer take to meet these environments MOST cost-effectively? (Select TWO)

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 52
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 54
A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

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

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

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

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

Answer: D

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

NEW QUESTION 55
A company receives food orders from multiple partners. The company has a microservices application that uses Amazon API Gateway APIs with AWS Lambda integration. Each partner sends orders by calling a customized API that is exposed through API Gateway. The API call invokes a shared Lambda function to process the orders.
Partners need to be notified after the Lambda function processes the orders. Each partner must receive updates for only the partner's own orders. The company wants to add new partners in the future with the fewest code changes possible.
Which solution will meet these requirements in the MOST scalable way?

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

Answer: C

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

NEW QUESTION 57
A developer is creating an application that will give users the ability to store photos from their cellphones in the cloud. The application needs to support tens of thousands of users. The application uses an Amazon API Gateway REST API that is integrated with AWS Lambda functions to process the photos. The application stores details about the photos in Amazon DynamoDB.
Users need to create an account to access the application. In the application, users must be able to upload photos and retrieve previously uploaded photos. The photos will range in size from 300 KB to 5 MB.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Cognito user pools to manage user account
B. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the AP
C. Use the Lambda function to store the photos and details in the DynamoDB tabl
D. Retrieve previously uploaded photos directly from the DynamoDB table.
E. Use Amazon Cognito user pools to manage user account
F. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the AP
G. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB tabl
H. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
I. Create an IAM user for each user of the application during the sign-up proces
J. Use IAM authentication to access the API Gateway AP

K. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB tabl
L. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
M. Create a users table in DynamoD
N. Use the table to manage user account
O. Create a Lambda authorizer that validates user credentials against the users tabl
P. Integrate the Lambda authorizer with API Gateway to control access to the AP
Q. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as par of the photo details in the DynamoDB tabl
R. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.

Answer: B

Explanation:

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

NEW QUESTION 62
A company has an application that runs as a series of AWS Lambda functions. Each Lambda function receives data from an Amazon Simple Notification Service (Amazon SNS) topic and writes the data to an Amazon Aurora DB instance.
To comply with an information security policy, the company must ensure that the Lambda functions all use a single securely encrypted database connection string to access Aurora.
Which solution will meet these requirements'?

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

Answer: A

Explanation:

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

NEW QUESTION 66
A developer must use multi-factor authentication (MFA) to access data in an Amazon S3
bucket that is in another AWS account. Which AWS Security Token Service (AWS STS) API operation should the developer use with the MFA information to meet this requirement?

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

Answer: D

Explanation:

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

NEW QUESTION 67
A company is building a web application on AWS. When a customer sends a request, the application will generate reports and then make the reports available to the customer within one hour. Reports should be accessible to the customer for 8 hours. Some reports are larger than 1 MB. Each report is unique to the customer. The application should delete all reports that are older than 2 days.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

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

NEW QUESTION 69
A company wants to share information with a third party. The third party has an HTTP API endpoint that the company can use to share the information. The company has the required API key to access the HTTP API.
The company needs a way to manage the API key by using code. The integration of the API key with the application code cannot affect application performance.
Which solution will meet these requirements MOST securely?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 72
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 75
A developer is creating a new REST API by using Amazon API Gateway and AWS Lambda. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.

The developer wants to make the REST API available for testing by using API Gateway locally.
Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

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

Answer: D

Explanation:
The AWS Serverless Application Model Command Line Interface (AWS SAM CLI) is a command-line tool for local development and testing of Serverless applications2. The sam local start-api subcommand of AWS SAM CLI is used to simulate a REST API by starting a new local endpoint3. Therefore, option D is correct.

NEW QUESTION 77
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions. When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment. If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

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

Answer: A

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

NEW QUESTION 80
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 84
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.

schedul
C. Create an Amazon EventBridge rule to invoke the Lambda function on a regular
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 88
A developer is trying get data from an Amazon DynamoDB table called demoman-table. The developer configured the AWS CLI to use a specific IAM use's credentials and ran the following command.

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

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

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

Answer: D

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

NEW QUESTION 93
A company built an online event platform For each event the company organizes quizzes and generates leaderboards that are based on the quiz scores. The company stores the leaderboard data in Amazon DynamoDB and retains the data for 30 days after an event is complete The company then uses a scheduled job to delete the old leaderboard data
The DynamoDB table is configured with a fixed write capacity. During the months when many events occur, the DynamoDB write API requests are throttled when the scheduled delete job runs.
A developer must create a long-term solution that deletes the old leaderboard data and optimizes write throughput
Which solution meets these requirements?

A. Configure a TTL attribute for the leaderboard data
B. Use DynamoDB Streams to schedule and delete the leaderboard data
C. Use AWS Step Functions to schedule and delete the leaderboard data.
D. Set a higher write capacity when the scheduled delete job runs

Answer: A

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

NEW QUESTION 98
A developer is building a serverless application by using AWS Serverless Application Model (AWS SAM) on multiple AWS Lambda functions.
When the application is deployed, the developer wants to shift 10% of the traffic to the new deployment of the application for the first 10 minutes after deployment.
If there are no issues, all traffic must switch over to the new version.
Which change to the AWS SAM template will meet these requirements?

A. Set the Deployment Preference Type to Canary10Percent10Minute
                                                                 AutoPublishAlias property to the Lambda alias.
B. Set the Deployment Preference Type to LinearIOPercentEvery10Minute
C. Set the
D. Set AutoPublishAlias property to the Lambda alias.
E. Set the Deployment Preference Type to CanaryIOPercentIOMinute
F. Set the PreTraffic and PostTraffic properties to the Lambda alias.
G. Set the Deployment Preference Type to LinearIOPercentEveryIOMinute
H. Set PreTraffic and Post Traffic properties to the Lambda alias.

Answer: A

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

NEW QUESTION 102
A company has an analytics application that uses an AWS Lambda function to process transaction data asynchronously A developer notices that asynchronous invocations of the Lambda function sometimes fail When failed Lambda function invocations occur, the developer wants to invoke a second Lambda function to handle errors and log details.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 107
A developer created an AWS Lambda function that performs a series of operations that involve multiple AWS services. The function's duration time is higher than normal. To determine the cause of the issue, the developer must investigate traffic between the services without changing the function code
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 108
A company has an ecommerce application. To track product reviews, the company's development team uses an Amazon DynamoDB table.
Every record includes the following
• A Review ID a 16-digrt universally unique identifier (UUID)
• A Product ID and User ID 16 digit UUIDs that reference other tables
• A Product Rating on a scale of 1-5
• An optional comment from the user
The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product.
Which index will provide the FASTEST response for this query"?

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

Answer: A

Explanation:
 This solution allows the fastest response for the query because it enables the query to use a single partition key value (the Product ID) and a range of sort key

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

NEW QUESTION 111
A developer is working on a web application that uses Amazon DynamoDB as its data store The application has two DynamoDB tables one table that is named artists and one table that is named songs The artists table has artistName as the partition key. The songs table has songName as the partition key and artistName as the sort key
The table usage patterns include the retrieval of multiple songs and artists in a single database operation from the webpage. The developer needs a way to retrieve this information with minimal network traffic and optimal application performance.
Which solution will meet these requirements'?

A. Perform a BatchGetItem operation that returns items from the two table
B. Use the list of songName artistName keys for the songs table and the list of artistName key for the artists table.
C. Create a local secondary index (LSI) on the songs table that uses artistName as the partition key Perform a query operation for each artistName on the songs table that filters by the list of songName Perform a query operation for each artistName on the artists table
D. Perform a BatchGetItem operation on the songs table that uses the songName/artistName key
E. Perform a BatchGetItem operation on the artists table that uses artistName as the key.
F. Perform a Scan operation on each table that filters by the list of songName/artistName for the songs table and the list of artistName in the artists table.

Answer: A

Explanation:
BatchGetItem can return one or multiple items from one or more tables. For reference check the link below
https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.h tml

NEW QUESTION 116
A developer uses AWS CloudFormation to deploy an Amazon API Gateway API and an AWS Step Functions state machine The state machine must reference the API Gateway API after the CloudFormation template is deployed The developer needs a solution that uses the state machine to reference the API Gateway endpoint.
Which solution will meet these requirements MOST cost-effectively?

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

Answer: A

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

NEW QUESTION 118
A company is running a custom application on a set of on-premises Linux servers that are accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled on the API test stage.
How can a developer enable X-Ray tracing on the on-premises servers with the LEAST amount of configuration?

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

Answer: B

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

References:
? [AWS X-Ray concepts - AWS X-Ray]
? [Setting up AWS X-Ray - AWS X-Ray]

NEW QUESTION 122
A developer has a legacy application that is hosted on-premises. Other applications hosted on AWS depend on the on-premises application for proper functioning.
In case of any application errors, the developer wants to be able to use Amazon CloudWatch to monitor and troubleshoot all applications from one place.
How can the developer accomplish this?

A. Install an AWS SDK on the on-premises server to automatically send logs to CloudWatch.
B. Download the CloudWatch agent to the on-premises serve
C. Configure the agent to use IAM user credentials with permissions for CloudWatch.
D. Upload log files from the on-premises server to Amazon S3 and have CloudWatch read the files.
E. Upload log files from the on-premises server to an Amazon EC2 instance and have the instance forward the logs to CloudWatch.

Answer: B

Explanation:
Amazon CloudWatch is a service that monitors AWS resources and applications. The developer can use CloudWatch to monitor and troubleshoot all applications from one place. To do so, the developer needs to download the CloudWatch agent to the on-premises server and configure the agent to use IAM user credentials with permissions for CloudWatch. The agent will collect logs and metrics from the on-premises server and send them to CloudWatch.
References:
? [What Is Amazon CloudWatch? - Amazon CloudWatch]
? [Installing and Configuring the CloudWatch Agent - Amazon CloudWatch]

NEW QUESTION 127
A developer warns to add request validation to a production environment Amazon API Gateway API. The developer needs to test the changes before the API is deployed to the production environment. For the lest the developer will send test requests to the API through a testing tool.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

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

NEW QUESTION 132
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 136
A developer is creating an AWS Lambda function in VPC mode An Amazon S3 event will invoke the Lambda function when an object is uploaded into an S3 bucket The Lambda function will process the object and produce some analytic results that will be recorded into a file Each processed object will also generate a log entry that will be recorded into a file.
Other Lambda functions. AWS services, and on-premises resources must have access to the result files and log file. Each log entry must also be appended to the same shared log file. The developer needs a solution that can share files and append results into an existing file.
Which solution should the developer use to meet these requirements?

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

Answer: A

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

NEW QUESTION 139
A developer is creating an Amazon DynamoDB table by using the AWS CLI The DynamoDB table must use server-side encryption with an AWS owned encryption key
How should the developer create the DynamoDB table to meet these requirements?

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

Answer: D

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

NEW QUESTION 142
A company is planning to use AWS CodeDeploy to deploy an application to Amazon Elastic Container Service (Amazon ECS) During the deployment of a new version of the application, the company initially must expose only 10% of live traffic to the new version of the deployed application. Then, after 15 minutes elapse, the company must route all the remaining live traffic to the new version of the deployed application.
Which CodeDeploy predefined configuration will meet these requirements?

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

Answer: A

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

NEW QUESTION 147
A developer is creating an AWS Lambda function. The Lambda function needs an external library to connect to a third-party solution The external library is a collection of files with a total size of 100 MB The developer needs to make the external library available to the Lambda execution environment and reduce the Lambda package space
Which solution will meet these requirements with the LEAST operational overhead?

A.

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

Answer: A

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

NEW QUESTION 149
An ecommerce application is running behind an Application Load Balancer. A developer observes some unexpected load on the application during non-peak hours. The developer wants to analyze patterns for the client IP addresses that use the application. Which HTTP header should the developer use for this analysis?

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

Answer: C

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

NEW QUESTION 153
A developer is investigating an issue in part of a company's application. In the application messages are sent to an Amazon Simple Queue Service (Amazon SQS) queue The AWS Lambda function polls messages from the SQS queue and sends email messages by using Amazon Simple Email Service (Amazon SES) Users have been receiving duplicate email messages during periods of high traffic.
Which reasons could explain the duplicate email messages? (Select TWO.)

A. Standard SQS queues support at-least-once message delivery
B. Standard SQS queues support exactly-once processing, so the duplicate email messages are because of user error.
C. Amazon SES has the DomainKeys Identified Mail (DKIM) authentication incorrectly configured
D. The SQS queue's visibility timeout is lower than or the same as the Lambda function's timeout.
E. The Amazon SES bounce rate metric is too high.

Answer: AD

Explanation:
 Standard SQS queues support at-least-once message delivery, which means that a message can be delivered more than once to the same or different

consumers. This can happen if the message is not deleted from the queue before the visibility timeout expires, or if there is a network issue or a system failure. The SQS queue's visibility timeout is the period of time that a message is invisible to other consumers after it is received by one consumer. If the visibility timeout is lower than or the same as the Lambda function's timeout, the Lambda function might not be able to process and delete the message before it becomes visible again, leading to duplicate processing and email messages. To avoid this, the visibility timeout should be set to at least 6 times the length of the Lambda function's timeout. The other options are not related to the issue of duplicate email messages. References

? Using the Amazon SQS message deduplication ID

? Exactly-once processing - Amazon Simple Queue Service

? Amazon SQS duplicated messages in queue - Stack Overflow

? amazon web services - How long can duplicate SQS messages persist …

? Standard SQS - Duplicate message | AWS re:Post - Amazon Web Services, Inc.

NEW QUESTION 158
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 162
A developer maintains applications that store several secrets in AWS Secrets Manager. The applications use secrets that have changed over time. The developer needs to identify required secrets that are still in use. The developer does not want to cause any application downtime.
What should the developer do to meet these requirements?

A. Configure an AWS CloudTrail log file delivery to an Amazon S3 bucke
B. Create an Amazon CloudWatch alarm for the GetSecretValu
C. Secrets Manager API operation requests
D. Create a secrets manager-secret-unused AWS Config managed rul
E. Create an Amazon EventBridge rule to Initiate notification when the AWS Config managed rule is met.
F. Deactivate the applications secrets and monitor the applications error logs temporarily.
G. Configure AWS X-Ray for the application
H. Create a sampling rule lo match the

GetSecretValue Secrets Manager API operation requests.

Answer: B

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

NEW QUESTION 163
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 164
A developer is writing a serverless application that requires an AWS Lambda function to be invoked every 10 minutes.
What is an automated and serverless way to invoke the function?

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

Answer: C

Explanation:
The solution that will meet the requirements is to create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function. This way, the developer can use an automated and serverless way to invoke the function every 10 minutes. The developer can also use a cron expression or a rate expression to specify the schedule for the rule. The other options either involve using an Amazon EC2 instance, which is not serverless, or using environment variables or query parameters, which do not trigger the function.
Reference: Schedule AWS Lambda functions using EventBridge

NEW QUESTION 168
A company is preparing to migrate an application to the company's first AWS environment Before this migration, a developer is creating a proof-of-concept application to validate a model for building and deploying container-based applications on AWS.
Which combination of steps should the developer take to deploy the containerized proof-of- concept application with the LEAST operational effort? (Select TWO.)

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 170
A developer wants to insert a record into an Amazon DynamoDB table as soon as a new file is added to an Amazon S3 bucket.
Which set of steps would be necessary to achieve this?

A. Create an event with Amazon EventBridge that will monitor the S3 bucket and then insert the records into DynamoDB.
B. Configure an S3 event to invoke an AWS Lambda function that inserts records into DynamoDB.
C. Create an AWS Lambda function that will poll the S3 bucket and then insert the records into DynamoDB.
D. Create a cron job that will run at a scheduled time and insert the records into DynamoDB.

Answer: B

Explanation:
Amazon S3 is a service that provides highly scalable, durable, and secure object storage. Amazon DynamoDB is a fully managed NoSQL database service that

provides fast and consistent performance with seamless scalability. AWS Lambda is a service that lets developers run code without provisioning or managing servers. The developer can configure an S3 event to invoke a Lambda function that inserts records into DynamoDB whenever a new file is added to the S3 bucket. This solution will meet the requirement of inserting a record into DynamoDB as soon as a new file is added to S3. References:
? [Amazon Simple Storage Service (S3)]
? [Amazon DynamoDB]
? [What Is AWS Lambda? - AWS Lambda]
? [Using AWS Lambda with Amazon S3 - AWS Lambda]

NEW QUESTION 175

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 178
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 182
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 186
A company wants to automate part of its deployment process. A developer needs to automate the process of checking for and deleting unused resources that supported previously deployed stacks but that are no longer used.
The company has a central application that uses the AWS Cloud Development Kit (AWS CDK) to manage all deployment stacks. The stacks are spread out across multiple accounts. The developer's solution must integrate as seamlessly as possible within the current deployment process.
Which solution will meet these requirements with the LEAST amount of configuration?

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

Answer: B

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

NEW QUESTION 187
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DVA-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DVA-C02 Product From:

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

Money Back Guarantee

DVA-C02 Practice Exam Features:

* DVA-C02 Questions and Answers Updated Frequently

* DVA-C02 Practice Questions Verified by Expert Senior Certified Staff

* DVA-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DVA-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year