Amazon

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

NEW QUESTION 1
A developer is creating an AWS Lambda function that needs credentials to connect to an Amazon RDS for MySQL database. An Amazon S3 bucket currently stores the credentials. The developer needs to improve the existing solution by implementing credential rotation and secure storage. The developer also needs to provide integration with the Lambda function.

Which solution should the developer use to store and retrieve the credentials with the LEAST management overhead?

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

Answer: C

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

NEW QUESTION 2
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 3

A developer is creating a simple proof-of-concept demo by using AWS CloudFormation and AWS Lambda functions The demo will use a CloudFormation template to deploy an existing Lambda function The Lambda function uses deployment packages and dependencies stored in Amazon S3 The developer defined anAWS Lambda Function resource in a CloudFormation template. The developer needs to add the S3 bucket to the CloudFormation template.
What should the developer do to meet these requirements with the LEAST development effort?

A. Add the function code in the CloudFormation template inline as the code property
B. Add the function code in the CloudFormation template as the ZipFile property.
C. Find the S3 key for the Lambda function Add the S3 key as the ZipFile property in the CloudFormation template.
D. Add the relevant key and bucket to the S3Bucket and S3Key properties in the CloudFormation template

Answer: D

Explanation:
 The easiest way to add the S3 bucket to the CloudFormation template is to use the S3Bucket and S3Key properties of the AWS::Lambda::Function resource. These properties specify the name of the S3 bucket and the location of the .zip file that contains the function code and dependencies. This way, the developer does not need to modify the function code or upload it to a different location. The other options are either not feasible or not efficient. The code property can only be used for inline code, not for code stored in S3. The ZipFile property can only be used for code that is less than 4096 bytes, not for code that has dependencies. Finding the S3 key for the Lambda function and adding it as the ZipFile property would not work, as the ZipFile property expects a base64-encoded .zip file, not an S3 location. References
? AWS::Lambda::Function - AWS CloudFormation
? Deploying Lambda functions as .zip file archives
? AWS Lambda Function Code - AWS CloudFormation

NEW QUESTION 4
A company has a web application that is hosted on Amazon EC2 instances The EC2 instances are configured to stream logs to Amazon CloudWatch Logs The company needs to receive an Amazon Simple Notification Service (Amazon SNS) notification when the number of application error messages exceeds a defined threshold within a 5-minute period
Which solution will meet these requirements?

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

Answer: D

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

NEW QUESTION 5
A developer is creating a mobile application that will not require users to log in. What is the MOST efficient method to grant users access to AWS resources'?

A. Use an identity provider to securely authenticate with the application.
B. Create an AWS Lambda function to create an 1AM user when a user accesses the application.
C. Create credentials using AWS KMS and apply these credentials to users when using the application.
D. Use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources.

Answer: D

Explanation:
 This solution is the most efficient method to grant users access to AWS resources without requiring them to log in. Amazon Cognito is a service that provides user sign-up, sign-in, and access control for web and mobile applications. Amazon Cognito identity pools support both authenticated and unauthenticated users. Unauthenticated users receive access to your AWS resources even if they aren't logged in with any of your identity providers (IdPs). You can use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources, such as Amazon S3 buckets or DynamoDB tables. This degree of access is useful to display content to users before they log in or to allow them to perform certain actions without signing up. Using an identity provider to securely authenticate with the application will require users to log in, which does not meet the requirement. Creating an AWS Lambda function to create an IAM user when a user accesses the application will incur unnecessary costs and complexity, and may pose security risks if not implemented properly. Creating credentials using AWS KMS and applying them to users when using the application will also incur unnecessary costs and complexity, and may not provide fine-grained access control for resources.
Reference: Switching unauthenticated users to authenticated users (identity pools), Allow
user access to your API without authentication (Anonymous user access)

NEW QUESTION 6
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 7
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 8
A company has an existing application that has hardcoded database credentials A developer needs to modify the existing application The application is deployed
in two AWS Regions with an active-passive failover configuration to meet company's disaster recovery strategy
The developer needs a solution to store the credentials outside the code. The solution must comply With the company's disaster recovery strategy
Which solution Will meet these requirements in the MOST secure way?

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

Answer: A

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

NEW QUESTION 9
A developer is working on an ecommerce platform that communicates with several third- party payment processing APIs The third-party payment services do not

provide a test environment.
The developer needs to validate the ecommerce platform's integration with the third-party payment processing APIs. The developer must test the API integration code without invoking the third-party payment processing APIs.
Which solution will meet these requirements'?

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

Answer: D

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

NEW QUESTION 10
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 10
A financial company must store original customer records for 10 years for legal reasons. A complete record contains personally identifiable information (PII). According to local regulations, PII is available to only certain people in the company and must not be shared with third parties. The company needs to make the records available to third-party organizations for statistical analysis without sharing the PII.
A developer wants to store the original immutable record in Amazon S3. Depending on who accesses the S3 document, the document should be returned as is or with all the PII removed. The developer has written an AWS Lambda function to remove the PII from the document. The function is named removePii.
What should the developer do so that the company can meet the PII requirements while maintaining only one copy of the document?

A. Set up an S3 event notification that invokes the removePii function when an S3 GET request is mad
B. Call Amazon S3 by using a GET request to access the object without PII.
C. Set up an S3 event notification that invokes the removePii function when an S3 PUT request is mad
D. Call Amazon S3 by using a PUT request to access the object without PII.
E. Create an S3 Object Lambda access point from the S3 consol
F. Select the removePii functio
G. Use S3 Access Points to access the object without PII.
H. Create an S3 access point from the S3 consol
I. Use the access point name to call the GetObjectLegalHold S3 API functio
J. Pass in the removePii function name to access the object without PII.

Answer: C

Explanation:
 S3 Object Lambda allows you to add your own code to process data retrieved from S3 before returning it to an application. You can use an AWS Lambda function to modify the data, such as removing PII, redacting confidential information, or resizing images. You can create an S3 Object Lambda access point and associate it

with your Lambda function. Then, you can use the access point to request objects from S3 and get the modified data back. This way, you can maintain only one copy of the original

document in S3 and apply different transformations depending on who accesses it. Reference: Using AWS Lambda with Amazon S3

NEW QUESTION 12
For a deployment using AWS Code Deploy, what is the run order of the hooks for in-place deployments?

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

Answer: B

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

NEW QUESTION 16
A developer has created an AWS Lambda function that is written in Python. The Lambda function reads data from objects in Amazon S3 and writes data to an Amazon DynamoDB table. The function is successfully invoked from an S3 event notification when an object is created. However, the function fails when it attempts to write to the DynamoDB table.
What is the MOST likely cause of this issue?

A. The Lambda function's concurrency limit has been exceeded.
B. DynamoDB table requires a global secondary index (GSI) to support writes.
C. The Lambda function does not have IAM permissions to write to DynamoDB.
D. The DynamoDB table is not running in the same Availability Zone as the Lambda function.

Answer: C

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_lambda- access-dynamodb.html

NEW QUESTION 20
A developer is testing an application that invokes an AWS Lambda function asynchronously. During the testing phase the Lambda function fails to process after two retries.
How can the developer troubleshoot the failure?

A. Configure AWS CloudTrail logging to investigate the invocation failures.
B. Configure Dead Letter Queues by sending events to Amazon SQS for investigation.
C. Configure Amazon Simple Workflow Service to process any direct unprocessed events.
D. Configure AWS Config to process any direct unprocessed events.

Answer: B

Explanation:
This solution allows the developer to troubleshoot the failure by capturing unprocessed events in a queue for further analysis. Dead Letter Queues (DLQs) are queues that store messages that could not be processed by a service, such as Lambda, for various reasons, such as configuration errors, throttling limits, or permissions issues. The developer can configure DLQs for Lambda functions by sending events to either an Amazon Simple Queue Service (SQS) queue or an Amazon Simple Notification Service (SNS) topic. The developer can then inspect the messages in the queue or topic to identify and fix the root cause of the failure. Configuring AWS CloudTrail logging will not capture invocation failures for asynchronous Lambda invocations, but only record API calls made by or on behalf of Lambda. Configuring Amazon Simple Workflow Service (SWF) or AWS Config will not process any direct unprocessed events, but require additional integration and configuration.
Reference: [Using AWS Lambda with DLQs], [Asynchronous invocation]

NEW QUESTION 21
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 24

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 28

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

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

Answer: C

Explanation:

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

NEW QUESTION 29

An application runs on multiple EC2 instances behind an ELB.
Where is the session data best written so that it can be served reliably across multiple requests?

A. Write data to Amazon ElastiCache
B. Write data to Amazon Elastic Block Store
C. Write data to Amazon EC2 instance Store

D. Wide data to the root filesystem

Answer: A

Explanation:
 The solution that will meet the requirements is to write data to Amazon ElastiCache. This way, the application can write session data to a fast, scalable, and reliable in-memory data store that can be served reliably across multiple requests. The other options either involve writing data to persistent storage, which is slower and more expensive than in-memory storage, or writing data to the root filesystem, which is not shared among multiple EC2 instances.
Reference: Using ElastiCache for session management

NEW QUESTION 30
A company is migrating its PostgreSQL database into the AWS Cloud. The company wants to use a database that will secure and regularly rotate database credentials. The company wants a solution that does not require additional programming overhead.
Which solution will meet these requirements?

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 33
A developer needs to store configuration variables for an application. The developer needs to set an expiration date and time for me configuration. The developer wants to receive notifications. Before the configuration expires. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:
 This solution will meet the requirements by creating an advanced parameter in AWS Systems Manager Parameter Store, which is a secure and scalable service for storing and managing configuration data and secrets. The advanced parameter allows setting expiration and expiration notification policy types, which enable specifying an expiration date and time for the configuration and receiving notifications before the configuration expires. The Lambda code will be refactored to load the Root CA Cert from the parameter store and modify the runtime trust store outside the Lambda function handler, which will improve performance and reduce latency by avoiding repeated calls to Parameter Store and trust store modifications for each invocation of the Lambda function. Option A is not optimal because it will create a standard parameter in AWS Systems Manager Parameter Store, which does not support expiration and expiration notification policy types. Option B is not optimal because it will create a secret access key and access key ID with permission to access the S3 bucket, which will introduce additional security risks and complexity for storing and managing credentials. Option D is not optimal because it will create a Docker container from Node.js base image to invoke Lambda functions, which will incur additional costs and overhead for creating and running Docker containers. References: AWS Systems Manager Parameter Store, [Using SSL/TLS to Encrypt a Connection to a DB Instance]

NEW QUESTION 35
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

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

Answer: B

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

NEW QUESTION 39
A developer is troubleshooting an application mat uses Amazon DynamoDB in the uswest- 2 Region. The application is deployed to an Amazon EC2 instance. The application requires read-only permissions to a table that is named Cars The EC2 instance has an attached IAM role that contains the following IAM policy.

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 41
A developer is preparing to begin development of a new version of an application. The previous version of the application is deployed in a production environment. The developer needs to deploy fixes and updates to the current version during the development of the new version of the application. The code for the new version of the application is stored in AWS CodeCommit.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 42
A development team maintains a web application by using a single AWS CloudFormation template. The template defines web servers and an Amazon RDS database. The team uses the Cloud Formation template to deploy the Cloud Formation stack to different environments.
During a recent application deployment, a developer caused the primary development database to be dropped and recreated. The result of this incident was a loss of data. The team needs to avoid accidental database deletion in the future.
Which solutions will meet these requirements? (Choose two.)

A. Add a CloudFormation Deletion Policy attribute with the Retain value to the database resource.
B. Update the CloudFormation stack policy to prevent updates to the database.
C. Modify the database to use a Multi-AZ deployment.
D. Create a CloudFormation stack set for the web application and database deployments.
E. Add a Cloud Formation DeletionPolicy attribute with the Retain value to the stack.

Answer: AB

Explanation:

AWS CloudFormation is a service that enables developers to model and provision AWS resources using templates. The developer can add a CloudFormation Deletion Policy attribute with the Retain value to the database resource. This will prevent the database from being deleted when the stack is deleted or updated. The developer can also update the CloudFormation stack policy to prevent updates to the database. This will prevent accidental changes to the database configuration or properties.
References:
? [What Is AWS CloudFormation? - AWS CloudFormation]
? [DeletionPolicy Attribute - AWS CloudFormation]
? [Protecting Resources During Stack Updates - AWS CloudFormation]

NEW QUESTION 46
A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function.
How should the developer configure the Lambda function to detect changes to the DynamoDB table?

A. Create an Amazon Kinesis data stream, and attach it to the DynamoDB tabl
B. Create a trigger to connect the data stream to the Lambda function.

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

Answer: C

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

NEW QUESTION 47
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 49
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
B. Set the AutoPublishAlias property to the Lambda alias.
C. Set the Deployment Preference Type to LinearlOPercentEvery10Minute
D. Set AutoPublishAlias property to the Lambda alias.
E. Set the Deployment Preference Type to CanarylOPercentIOMinute
F. Set the PreTraffic and PostTraffic properties to the Lambda alias.
G. Set the Deployment Preference Type to LinearlOPercentEveryIOMinute
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 51
An application that runs on AWS Lambda requires access to specific highly confidential objects in an Amazon S3 bucket. In accordance with the principle of least

privilege a company grants access to the S3 bucket by using only temporary credentials.
How can a developer configure access to the S3 bucket in the MOST secure way?

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

Answer: C

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

NEW QUESTION 54

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

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

Answer: BD

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

NEW QUESTION 58
A developer is deploying an AWS Lambda function The developer wants the ability to return to older versions of the function quickly and seamlessly.
How can the developer achieve this goal with the LEAST operational overhead?

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

Answer: B

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

NEW QUESTION 59
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 60
A company's developer has deployed an application in AWS by using AWS CloudFormation The CloudFormation stack includes parameters in AWS Systems Manager Parameter Store that the application uses as configuration settings. The application can modify the parameter values
When the developer updated the stack to create additional resources with tags, the developer noted that the parameter values were reset and that the values ignored the latest changes made by the application. The developer needs to change the way the company deploys the CloudFormation stack. The developer also needs to avoid resetting the parameter values outside the stack.
Which solution will meet these requirements with the LEAST development effort?

A. Modify the CloudFormation stack to set the deletion policy to Retain for the Parameter Store parameters.
B. Create an Amazon DynamoDB table as a resource in the CloudFormation stack to hold configuration data for the application Migrate the parameters that the application is modifying from Parameter Store to the DynamoDB table
C. Create an Amazon RDS DB instance as a resource in the CloudFormation stac
D. Create a table in the database for parameter configuratio
E. Migrate the parameters that the application is modifying from Parameter Store to the configuration table
F. Modify the CloudFormation stack policy to deny updates on Parameter Store parameters

Answer: D

Explanation:
 https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack- resources.html#stack-policy-samples

NEW QUESTION 64
A developer is using AWS Amplify Hosting to build and deploy an application. The developer is receiving an increased number of bug reports from users. The developer wants to add end-to-end testing to the application to eliminate as many bugs as possible before the bugs reach production.
Which solution should the developer implement to meet these requirements?

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

Answer: C

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

NEW QUESTION 69
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 74
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 76
A developer deployed an application to an Amazon EC2 instance The application needs to know the public IPv4 address of the instance
How can the application find this information?

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

Answer: A

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

NEW QUESTION 77
A team of developed is using an AWS CodePipeline pipeline as a continuous integration and continuous delivery (CI/CD) mechanism for a web application. A developer has written unit tests to programmatically test the functionality of the application code. The unit tests produce a test report that shows the results of each individual check. The developer now
wants to run these tests automatically during the CI/CD process.

A. Write a Git pre-commit hook that runs the test before every commi
B. Ensure that each developer who is working on the project has the pre-commit hook instated locall
C. Review the test report and resolve any issues before pushing changes to AWS CodeCommit.
D. Add a new stage to the pipelin
E. Use AWS CodeBuild as the provide
F. Add the new stage after the stage that deploys code revisions to the test environmen
G. Write a buildspec that fails the CodeBuild stage if any test does not pas
H. Use the test reports feature of Codebuild to integrate the report with the CodoBuild consol
I. View the test results in CodeBuild Resolve any issues.
J. Add a new stage to the pipelin
K. Use AWS CodeBuild at the provide
L. Add the new stage before the stage that deploys code revisions to the test environmen
M. Write a buildspec that fails the CodeBuild stage it any test does not pas
N. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild consol
O. View the test results in codeBuild Resolve any issues.
P. Add a new stage to the pipelin
Q. Use Jenkins as the provide
R. Configure CodePipeline to use Jenkins to run the unit test
S. Write a Jenkinsfile that fails the stage if any test does not pas
T. Use the test report plugin for Jenkins to integrate the repot with the Jenkins dashboar
. View the test results in Jenkin
. Resolve any issues.

Answer: C

Explanation:
The solution that will meet the requirements is to add a new stage to the pipeline. Use AWS CodeBuild as the provider. Add the new stage before the stage that deploys code revisions to the test environment. Write a buildspec that fails the CodeBuild stage if any test does not pass. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild console. View the test results in CodeBuild. Resolve any issues. This way, the developer can run the unit tests automatically during the CI/CD process and catch any bugs before deploying to the test environment. The developer can also use the test reports feature of CodeBuild to view and analyze the test results in a graphical interface. The other options either involve running the tests manually, running them after deployment, or using a different provider that requires additional configuration and integration.
Reference: Test reports for CodeBuild

NEW QUESTION 80
A developer is planning to migrate on-premises company data to Amazon S3. The data must be encrypted, and the encryption Keys

must support automate annual rotation. The company must use AWS Key Management Service (AWS KMS) to encrypt the data.
When type of keys should the developer use to meet these requirements?

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

Answer: B

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

NEW QUESTION 85
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 86
A developer is working on an ecommerce website The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available
How can the developer update the application to meet these requirements with MINIMUM changes?

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

Answer: D

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

NEW QUESTION 89
An application is using Amazon Cognito user pools and identity pools for secure access. A developer wants to integrate the user-specific file upload and download features in the application with Amazon S3. The developer must ensure that the files are saved and retrieved in a secure manner and that users can access only their own files. The file sizes range from 3 KB to 300 MB.
Which option will meet these requirements with the HIGHEST level of security?

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

Answer: D

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

NEW QUESTION 90
A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment, the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.
Which deployment method should the developer use to meet these requirements?

A.

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

Answer: D

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

NEW QUESTION 93
A developer is building an application that gives users the ability to view bank account from multiple sources in a single dashboard. The developer has automated the process to retrieve API credentials for these sources. The process invokes an AWS Lambda function that is associated with an AWS CloudFormation cotton resource.
The developer wants a solution that will store the API credentials with minimal operational overhead.
When solution will meet these requirements?

A. Add an AWS Secrets Manager GenerateSecretString resource to the CloudFormation templat
B. Set the value to reference new credentials to the Cloudformation resource.
C. Use the AWS SDK ssm PutParameter operation in the Lambda function from the existing, custom resource to store the credentials as a paramete
D. Set the parameter value to reference the new credential
E. Set ma parameter type to SecureString.
F. Add an AWS Systems Manager Parameter Store resource to the CloudFormation templat
G. Set the CloudFormation resource value to reference the new credentials Set the resource NoEcho attribute to true.
H. Use the AWS SDK ssm PutParameter operation in the Lambda function from the existing custom resources to store the credentials as a paramete
I. Set the parameter value to reference the new credential
J. Set the parameter NoEcho attribute to true.

Answer: B

Explanation:
The solution that will meet the requirements is to use the AWS SDK ssm PutParameter operation in the Lambda function from the existing custom resource to store the credentials as a parameter. Set the parameter value to reference the new credentials. Set the parameter type to SecureString. This way, the developer can store the API credentials with minimal operational overhead, as AWS Systems Manager Parameter Store provides secure and scalable storage for configuration data. The SecureString parameter type encrypts the parameter value with AWS Key Management Service (AWS KMS). The other options either involve adding additional resources to the CloudFormation template, which increases complexity and cost, or do not encrypt the parameter value, which reduces security.
Reference: Creating Systems Manager parameters

NEW QUESTION 95
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 99
An ecommerce company is using an AWS Lambda function behind Amazon API Gateway

as its application tier. To process orders during checkout, the application calls a POST API from the frontend. The POST API invokes the Lambda function asynchronously. In rare situations, the application has not processed orders. The Lambda application logs show no errors or failures.
What should a developer do to solve this problem?

A. Inspect the frontend logs for API failure
B. Call the POST API manually by using the requests from the log file.
C. Create and inspect the Lambda dead-letter queu
D. Troubleshoot the failed function
E. Reprocess the events.
F. Inspect the Lambda logs in Amazon CloudWatch for possible error
G. Fix the errors.
H. Make sure that caching is disabled for the POST API in API Gateway.

Answer: B

Explanation:
The solution that will solve this problem is to create and inspect the Lambda dead-letter queue. Troubleshoot the failed functions. Reprocess the events. This way, the developer can identify and fix any issues that caused the Lambda function to fail when invoked asynchronously by API Gateway. The developer can also reprocess any orders that were not processed due to failures. The other options either do not address the root cause of the problem, or do not help recover from

failures.
Reference: Asynchronous invocation

NEW QUESTION 102
A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the AWS Key Management Service (AWS KMS) key was used and by whom.
Which encryption option will meet these requirements?

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

Answer: B

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

NEW QUESTION 105
A developer is incorporating AWS X-Ray into an application that handles personal identifiable information (PII). The application is hosted on Amazon EC2 instances. The application trace messages include encrypted PII and go to Amazon CloudWatch. The developer needs to ensure that no PII goes outside of the EC2 instances.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 110
A developer is building an application that uses AWS API Gateway APIs. AWS Lambda function, and AWS Dynamic DB tables. The developer uses the AWS Serverless Application Model (AWS SAM) to build and run serverless applications on AWS. Each time the developer pushes of changes for only to the Lambda functions, all the artifacts in the application are rebuilt.
The developer wants to implement AWS SAM Accelerate by running a command to only redeploy the Lambda functions that have changed.
Which command will meet these requirements?

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

Answer: D

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

NEW QUESTION 113
A developer is creating an application that includes an Amazon API Gateway REST API in the us-east-2 Region. The developer wants to use Amazon CloudFront and a custom domain name for the API. The developer has acquired an SSL/TLS certificate for the domain from a third-party provider.
How should the developer configure the custom domain for the application?

A. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the AP
B. Create a DNS A record for the custom domain.

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

Answer: D

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

NEW QUESTION 117
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 122
A developer needs to deploy an application running on AWS Fargate using Amazon ECS The application has environment variables that must be passed to a container for the application to initialize.
How should the environment variables be passed to the container?

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

Answer: B

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

NEW QUESTION 124
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 127
A developer is testing a RESTful application that is deployed by using Amazon API Gateway and AWS Lambda When the developer tests the user login by using credentials that are not valid, the developer receives an HTTP 405 METHOD_NOT_ALLOWED error The developer has verified that the test is sending the correct request for the resource
Which HTTP error should the application return in response to the request?

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

Answer: A

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

NEW QUESTION 132
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 137
A developer is optimizing an AWS Lambda function and wants to test the changes in
                         production on a small percentage of all traffic. The Lambda function serves requests to a REST API in Amazon API Gateway. The developer needs to deploy their changes and perform a test in production without changing the API Gateway URL.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 140
A developer is designing a serverless application for a game in which users register and log in through a web browser The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API
The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.
Which solution will meet these requirements'?

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

Answer: A

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

NEW QUESTION 144
A developer needs to build an AWS CloudFormation template that self-populates the AWS Region variable that deploys the CloudFormation template
What is the MOST operationally efficient way to determine the Region in which the template is being deployed?

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

Answer: A

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

NEW QUESTION 145
A company is building a serverless application on AWS. The application uses an AWS Lambda function to process customer orders 24 hours a day, 7 days a week. The Lambda function calls an external vendor's HTTP API to process payments.
During load tests, a developer discovers that the external vendor payment processing API occasionally times out and returns errors. The company expects that some payment processing API calls will return errors.
The company wants the support team to receive notifications in near real time only when the payment processing external API error rate exceed 5% of the total number of transactions in an hour. Developers need to use an existing Amazon Simple Notification Service (Amazon SNS) topic that is configured to notify the support team.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 149
An Amazon Kinesis Data Firehose delivery stream is receiving customer data that contains personally identifiable information. A developer needs to remove pattern-based customer identifiers from the data and store the modified data in an Amazon S3 bucket.
What should the developer do to meet these requirements?

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

Answer: A

Explanation:
 Amazon Kinesis Data Firehose is a service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon OpenSearch Service, and Amazon Kinesis Data Analytics. The developer can implement Kinesis Data Firehose data transformation as an AWS Lambda function. The function can remove pattern-based customer identifiers from the data and return the modified data to Kinesis Data Firehose. The developer can set an Amazon S3 bucket as the destination of the delivery stream. References:
? [What Is Amazon Kinesis Data Firehose? - Amazon Kinesis Data Firehose]
? [Data Transformation - Amazon Kinesis Data Firehose]

NEW QUESTION 152

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

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

Answer: A

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

NEW QUESTION 156

A company hosts its application on AWS. The application runs on an Amazon Elastic Container Service (Amazon ECS) cluster that uses AWS Fargate. The cluster runs behind an Application Load Balancer The application stores data in an Amazon Aurora database A developer encrypts and manages database credentials inside the application

The company wants to use a more secure credential storage method and implement periodic credential rotation.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

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

? Top 5 AWS root account best practices - TheServerSide

NEW QUESTION 160
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 has an Amazon S3 bucket that contains sensitive data. The data must be encrypted in transit and at rest. The company encrypts the data in the S3 bucket by using an AWS Key Management Service (AWS KMS) key. A developer needs to grant several other AWS accounts the permission to use the S3 GetObject operation to retrieve the data from the S3 bucket.

How can the developer enforce that all requests to retrieve the data provide encryption in transit?

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

Answer: A

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

NEW QUESTION 164
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 169

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

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

Answer: D

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

NEW QUESTION 173

......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Certified-Developer-Associate Practice Exam Features:

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

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

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

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

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