Free Questions for DVA-C02 by certsinside

Shared by Wilkerson on 09-08-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company is using Amazon API Gateway to invoke a new AWS Lambda function The company has Lambda function versions in its PROD and DEV environments. In each environment, there is a Lambda function alias pointing to the corresponding Lambda function version API Gateway has one stage that is configured to point at the PROD alias

The company wants to configure API Gateway to enable the PROD and DEV Lambda function versions to be simultaneously and distinctly available

Which solution will meet these requirements?

Options:

A- Enable a Lambda authorizer for the Lambda function alias in API Gateway Republish PROD and create a new stage for DEV Create API Gateway stage variables for the PROD and DEV stages. Point each stage variable to the PROD Lambda authorizer to the DEV Lambda authorizer.

B- Set up a gateway response in API Gateway for the Lambda function alias. Republish PROD and create a new stage for DEV. Create gateway responses in API Gateway for PROD and DEV Lambda aliases

C- Use an environment variable for the Lambda function alias in API Gateway. Republish PROD and create a new stage for development. Create API gateway environment variables for PROD and DEV stages. Point each stage variable to the PROD Lambda function alias to the DEV Lambda function alias.

D- Use an API Gateway stage variable to configure the Lambda function alias Republish PROD and create a new stage for development Create API Gateway stage variables for PROD and DEV stages Point each stage variable to the PROD Lambda function alias and to the DEV Lambda function alias

Answer:

D

Explanation:

API Gateway Stages:Stages in API Gateway represent distinct environments (like PROD and DEV) allowing different configurations.

Stage Variables:Stage variables store environment-specific information, including Lambda function aliases.

Ease of Management:This solution offers a straightforward way to manage different Lambda function versions across environments.

API Gateway Stages:https://docs.aws.amazon.com/apigateway/latest/developerguide/set-up-stages.html

API Gateway Stage Variables:https://docs.aws.amazon.com/apigateway/latest/developerguide/stage-variables.html

Question 2

Question Type: MultipleChoice

A company runs a payment application on Amazon EC2 instances behind an Application Load Balance The EC2 instances run in an Auto Scaling group across multiple Availability Zones The application needs to retrieve application secrets during the application startup and export the secrets as environment variables These secrets must be encrypted at rest and need to be rotated every month.

Which solution will meet these requirements with the LEAST development effort?

Options:

A- Save the secrets in a text file and store the text file in Amazon S3 Provision a customer managed key Use the key for secret encryption in Amazon S3 Read the contents of the text file and read the export as environment variables Configure S3 Object Lambda to rotate the text file every month

B- Save the secrets as strings in AWS Systems Manager Parameter Store and use the default AWS Key Management Service (AWS KMS) key Configure an Amazon EC2 user data script to retrieve the secrets during the startup and export as environment variables Configure an AWS Lambda function to rotate the secrets in Parameter Store every month.

C- Save the secrets as base64 encoded environment variables in the application properties. Retrieve the secrets during the application startup. Reference the secrets in the application code. Write a script to rotate the secrets saved as environment variables.

D- Store the secrets in AWS Secrets Manager Provision a new customer master key Use the key to encrypt the secrets Enable automatic rotation Configure an Amazon EC2 user data script to programmatically retrieve the secrets during the startup and export as environment variables

Answer:

D

Explanation:

AWS Secrets Manager:Built for managing secrets, providing encryption, automatic rotation, and access control.

Customer Master Key (CMK):Provides an extra layer of control over encryption through AWS KMS.

Automatic Rotation:Enhances security by regularly changing the secret.

User Data Script:Allows secrets retrieval at instance startup and sets them as environment variables for seamless use within the application.

AWS Secrets Manager Documentation:https://docs.aws.amazon.com/secretsmanager/

AWS KMS Documentation:https://docs.aws.amazon.com/kms/

User Data for EC2 Instances:https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html

Question 3

Question Type: MultipleChoice

A company has a social media application that receives large amounts of traffic User posts and interactions are continuously updated in an Amazon RDS database The data changes frequently, and the data types can be complex The application must serve read requests

with minimal latency

The application's current architecture struggles to deliver these rapid data updates efficiently The company needs a solution to improve the application's performance.

Which solution will meet these requirements'?

Options:

A- Use Amazon DynamoDB Accelerator (DAX) in front of the RDS database to provide a caching layer for the high volume of rapidly changing data

B- Set up Amazon S3 Transfer Acceleration on the RDS database to enhance the speed of data transfer from the databases to the application.

C- Add an Amazon CloudFront distribution in front of the RDS database to provide a caching layer for the high volume of rapidly changing data

D- Create an Amazon ElastiCache for Redis cluster. Update the application code to use a write-through caching strategy and read the data from Redis.

Answer:

D

Explanation:

Amazon ElastiCache for Redis:An in-memory data store known for extremely low latency, ideal for caching frequently accessed, complex data.

Write-Through Caching:Ensures that data is always consistent between the cache and the database. Writes go to both Redis and RDS.

Performance Gains:Redis handles reads with minimal latency, offloading the RDS database and improving the application's responsiveness.

Amazon ElastiCache for Redis Documentation:https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/

Caching Strategies:https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Strategies.html

Question 4

Question Type: **MultipleChoice**

A developer is building a microservices-based application by using Python on AWS and several AWS services The developer must use AWS X-Ray The developer views the service map by using the console to view the service dependencies. During testing, the developer notices that some services are missing from the service map

What can the developer do to ensure that all services appear in the X-Ray service map?

Options:

A- Modify the X-Ray Python agent configuration in each service to increase the sampling rate

B- Instrument the application by using the X-Ray SDK for Python. Install the X-Ray SDK for all the services that the application uses

C- Enable X-Ray data aggregation in Amazon CloudWatch Logs for all the services that the application uses

D- Increase the X-Ray service map timeout value in the X-Ray console

Answer:

B

Explanation:

AWS X-Ray SDK:The primary way to enable X-Ray tracing within applications. The SDK sends data about requests and subsegments to the X-Ray daemon for service map generation.

Instrumenting All Services:To visualize a complete microservice architecture on the service map, each relevant service must include the X-Ray SDK.

AWS X-Ray Documentation:https://docs.aws.amazon.com/xray/

X-Ray SDK for Python:https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-python.html

Question 5

Question Type: MultipleChoice

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

Which solution will meet these requirements?

Options:

A- Create an Amazon Simple Notification Service (Amazon SNS) topic Configure S3 event notifications with a destination of the SNS topic Subscribe the Lambda function to the SNS topic Create an email notification subscription to the SNS topic

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

C- Create an Amazon Simple Queue Service (Amazon SQS) queue Configure S3 event notifications with a destination of the SQS queue Subscribe the Lambda function to the SQS queue Create an email notification subscription to the SQS queue.

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

Answer:

A

Explanation:

SNS as a Fan-out Mechanism:SNS is perfect for triggering multiple actions from a single event (here, the image upload).

Workflow:

SNS Topic:Create an SNS topic that will be the central notification point.

S3 Event Notification:Configure the S3 bucket to send 'Object Created' event notifications to the SNS topic.

Lambda Subscription:Subscribe your thumbnail-creating Lambda function to the SNS topic.

Email Subscription:Subscribe an email address to the SNS topic to trigger notifications.

S3 Event Notifications:https://docs.aws.amazon.com/AmazonS3/latest/userguide/EventNotifications.html

SNS Subscriptions:https://docs.aws.amazon.com/sns/latest/dg/SNSMobilePush.html

Question 6

Question Type: MultipleChoice

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?

Options:

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 group. Configure the filter to send an SNS notification when the number of errors exceeds the defined threshold within a 5-minute period.

C- 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

D- 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 metric. 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:

CloudWatch for Log Analysis:CloudWatch is the best fit here because logs are already centralized. Here's the process:

Metric Filter: Create a metric filter on the CloudWatch Logs log group. Design a pattern to specifically identify application error messages.

Custom Metric: This filter generates a new custom CloudWatch metric (e.g.,ApplicationErrors). This metric tracks the error count.

CloudWatch Alarm: Create an alarm on theApplicationErrorsmetric. Configure the alarm with your desired threshold and a 5-minute evaluation period.

SNS Action: Set the alarm to trigger an SNS notification when it enters the alarm state.

CloudWatch Metric Filters:https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html

CloudWatch Alarms:https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html

Question 7

Question Type: MultipleChoice

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?

Options:

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:

Instance Metadata Service:EC2 instances have access to an internal metadata service. It provides instance-specific information like instance ID, security groups, and public IP address.

Accessing Metadata:

Make an HTTP GET request to the base URL:http://169.254.169.254/latest/meta-data/

You'll get a list of available categories. The public IPv4 address is underpublic-ipv4.

Instance Metadata and User Data:https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html

Question 8

Question Type: MultipleChoice

A company runs an application on AWS The application uses an AWS Lambda function that is configured with an Amazon Simple Queue Service (Amazon SQS) queue called high priority queue as the event source A developer is updating the Lambda function with another SQS queue called low priority queue as the event source The Lambda function must always read up to 10 simultaneous messages from the high priority queue before processing messages from low priority queue. The Lambda function must be limited to 100 simultaneous invocations.

Which solution will meet these requirements'?

Options:

A- Set the event source mapping batch size to 10 for the high priority queue and to 90 for the low priority queue

B- Set the delivery delay to 0 seconds for the high priority queue and to 10 seconds for the low priority queue

C- Set the event source mapping maximum concurrency to 10 for the high priority queue and to 90 for the low priority queue

D- Set the event source mapping batch window to 10 for the high priority queue and to 90 for the low priority queue

Answer:

C

Explanation:

Lambda Concurrency:The 'maximum concurrency' setting in event source mappings controls the maximum number of simultaneous invocations Lambda allows for that specific source.

Prioritizing Queues:Setting a lower maximum concurrency for the 'high priority queue' ensures it's processed first while allowing more concurrent invocations from the 'low priority queue'.

Batching:Batch size settings affect the number of messages Lambda retrieves from a queue per invocation, which is less relevant to the prioritization requirement.

Lambda Event Source Mappings:https://docs.aws.amazon.com/lambda/latest/dg/invocation-eventsourcemapping.html

Lambda Concurrency:https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html

Question 9

Question Type: MultipleChoice

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?

Options:

A- HTTP 401

B- HTTP 404

C- HTTP 503

D- HTTP 505

Answer:

A

Explanation:

HTTP Status Codes:Each HTTP status code has a specific meaning in RESTful APIs.

HTTP 405 (Method Not Allowed):Indicates that the request method (e.g., POST) is not supported for the specified resource.

HTTP 401 (Unauthorized):Represents a failure to authenticate, which is the appropriate response for invalid login credentials.

HTTP Status Codes:https://developer.mozilla.org/en-US/docs/Web/HTTP/Status