Free Questions for SAP-C02 by ebraindumps

Shared by Stuart on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company is deploying AWS Lambda functions that access an Amazon RDS for PostgreSQL database. The company needs to launch the Lambda functions in a QA

environment and in a production environment.

The company must not expose credentials within application code and must rotate passwords automatically.

Which solution will meet these requirements?

Options:

A- Store the database credentials for both environments in AWS Systems Manager Parameter Store. Encrypt the credentials by using an AWS Key Management
Service (AWS KMS) key. Within the application code of the Lambda functions, pull the credentials from the Parameter Store parameter by using the AWS SDK
for Python (Bot03). Add a role to the Lambda functions to provide access to the Parameter Store parameter.

B- Store the database credentials for both environments in AWS Secrets Manager with distinct key entry for the QA environment and the production environment.
Turn on rotation. Provide a reference to the Secrets Manager key as an environment variable for the Lambda functions.

C- Store the database credentials for both environments in AWS Key Management Service (AWS KMS). Turn on rotation. Provide a reference to the credentials
that are stored in AWS KMS as an environment variable for the Lambda functions.

D- Create separate S3 buckets for the QA environment and the production environment. Turn on server-side encryption with AWS KMS keys (SSE-KMS) for the
S3 buckets. Use an object naming pattern that gives each Lambda function's application code the ability to pull the correct credentials for the function's
corresponding environment. Grant each Lambda function's execution role access to Amazon S3.

Answer:

B

Explanation:

The best solution is to store the database credentials for both environments in AWS Secrets Manager with distinct key entry for the QA environment and the production environment. AWS Secrets Manager is a web service that can securely store, manage, and retrieve secrets, such as database credentials. AWS Secrets Manager also supports automatic rotation of secrets by using Lambda functions or built-in rotation templates. By storing the database credentials for both environments in AWS Secrets Manager, the company can avoid exposing credentials within application code and rotate passwords automatically. By providing a reference to the Secrets Manager key as an environment variable for the Lambda functions, the company can easily access the credentials from the code by using the AWS SDK. This solution meets all the requirements of the company.

Question 2

Question Type: MultipleChoice

A company has mounted sensors to collect information about environmental parameters such as humidity and light throughout all the company's factories. The company needs to stream and analyze the data in the AWS Cloud in real time. If any of the parameters fall out of acceptable ranges, the factory operations team must receive a notification immediately.

Which solution will meet these requirements?

Options:

A- Stream the data to an Amazon Kinesis Data Firehose delivery stream. Use AWS Step Functions to consume and analyze the data in the Kinesis Data Firehose
delivery stream. use Amazon Simple Notification Service (Amazon SNS) to notify the operations team.

B- Stream the data to an Amazon Managed Streaming for Apache Kafka (Amazon MSK) cluster. Set up a trigger in Amazon MSK to invoke an AWS Fargate task
to analyze the data. Use Amazon Simple Email Service (Amazon SES) to notify the operations team.

C- Stream the data to an Amazon Kinesis data stream. Create an AWS Lambda function to consume the Kinesis data stream and to analyze the data. Use
Amazon Simple Notification Service (Amazon SNS) to notify the operations team.

D- Stream the data to an Amazon Kinesis Data Analytics application. I-Jse an automatically scaled and containerized service in Amazon Elastic Container Service (Amazon ECS) to consume and analyze the data. use Amazon Simple Email Service (Amazon SES) to notify the operations team.

Answer:

C

Explanation:

The best solution is to stream the data to an Amazon Kinesis data stream and create an AWS Lambda function to consume the Kinesis data stream and to analyze the data. Amazon Kinesis is a web service that can collect, process, and analyze real-time streaming data from various sources, such as sensors. AWS Lambda is a serverless computing service that can run code in response to events, such as incoming data from a Kinesis data stream. By using AWS Lambda, the company can avoid provisioning or managing servers and scale automatically based on the demand. Amazon Simple Notification Service (Amazon SNS) is a web service that enables applications to send and receive notifications from the cloud. By using Amazon SNS, the company can notify the operations team immediately if any of the parameters fall out of acceptable ranges. This solution meets all the requirements of the company.

Question 3

Question Type: MultipleChoice

A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time. A user is notified when image processing is complete.

Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load? (Select THREE.)

Options:

A- Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.

B- Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SOS) standard queue.

C- Invoke an AWS Lambda function to perform image processing when a message is available in the queue.

D- Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue

E- Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete.

F- Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

Answer:

B, C, E

Explanation:

The best solution is to upload files from the mobile software directly to Amazon S3, use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue, and invoke an AWS Lambda function to perform image processing when a message is available in the queue. This solution will ensure that image processing can scale to handle the load, as Amazon S3 can store any amount of data and handle concurrent uploads, Amazon SQS can buffer the messages and deliver them reliably, and AWS Lambda can run code without provisioning or managing servers and scale automatically based on the demand. This solution will also notify the user when processing is complete by sending a push notification to the mobile app using Amazon Simple Notification Service (Amazon SNS), which is a web service that enables applications to send and receive notifications from the cloud. This solution is more cost-effective than using Amazon MQ, which is a managed message broker service for Apache ActiveMQ that requires a dedicated broker instance, or S3 Batch Operations, which is a feature that allows users to perform bulk actions on S3 objects, such as copying or tagging, but does not support custom code execution. This solution is also more suitable than using Amazon Simple Email Service (Amazon SES), which is a web service that enables applications to send and receive email messages, but does not support push notifications for mobile devices.Reference:Amazon S3 Documentation,Amazon SQS Documentation,AWS Lambda Documentation,Amazon SNS Documentation

Question 4

Question Type: MultipleChoice

An online survey company runs its application in the AWS Cloud. The application is distributed and consists of microservices that run in an automatically scaled Amazon Elastic Container Service (Amazon ECS) cluster. The ECS cluster is a target for an Application Load

Balancer (ALB). The ALB is a custom origin for an Amazon CloudFront distribution.

The company has a survey that contains sensitive dat

a. The sensitive data must be encrypted when it moves through the application. The application's data-handling microservice is the only microservice that should be able to decrypt the data.

Which solution will meet these requirements?

Options:

A- Create a symmetric AWS Key Management Service (AWS KMS) key that is dedicated to the data-handling microservice. Create a field-level encryption profile and a configuration. Associate the KMS key and the configuration with the CloudFront cache behavior.

B- Create an RSA key pair that is dedicated to the data-handling microservice. Upload the public key to the CloudFront distribution. Create a field-level encryption profile and a configuration. Add the configuration to the CloudFront cache behavior.

C- Create a symmetric AWS Key Management Service (AWS KMS) key that is dedicated to the data-handling microservice. Create a Lambda@Edge function. Program the function to use the KMS key to encrypt the sensitive data.

D- Create an RSA key pair that is dedicated to the data-handling microservice. Create a Lambda@Edge function. Program the function to use the private key of the RSA key pair to encrypt the sensitive data.

Answer:

B

Explanation:

The best solution is to create an RSA key pair that is dedicated to the data-handling microservice and upload the public key to the CloudFront distribution. Then, create a field-level encryption profile and a configuration, and add the configuration to the CloudFront cache behavior. This solution will ensure that the sensitive data is encrypted at the edge locations of CloudFront, close to the end users, and remains encrypted throughout the application stack. Only the data-handling microservice, which has access to the private key of the RSA key pair, can decrypt the data. This solution does not require any additional resources or code changes, and leverages the built-in feature of CloudFront field-level encryption. For more information about CloudFront field-level encryption, seeUsing field-level encryption to help protect sensitive data.

Question 5

Question Type: MultipleChoice

A company uses an organization in AWS Organizations to manage the company's AWS accounts. The company uses AWS CloudFormation to deploy all infrastructure. A finance team wants to buikJ a chargeback model The finance team asked each business unit to tag resources by using a predefined list of project values.

When the finance team used the AWS Cost and Usage Report in AWS Cost Explorer and filtered based on project, the team noticed noncompliant project values. The company wants to enforce the use of project tags for new resources.

Which solution will meet these requirements with the LEAST effort?

Options:

A- Create a tag policy that contains the allowed project tag values in the organization's management account. Create an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added. Attach the SCP to each OU.

B- Create a tag policy that contains the allowed project tag values in each OU. Create an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added. Attach the SCP to each OU.

C- Create a tag policy that contains the allowed project tag values in the AWS management account. Create an 1AM policy that denies the cloudformation:CreateStack API operation unless a project tag is added. Assign the policy to each user.

D- Use AWS Service Catalog to manage the CloudFoanation stacks as products. Use a TagOptions library to control project tag values. Share the portfolio with all OUs that are in the organization.

Answer:

A

Explanation:

The best solution is to create a tag policy that contains the allowed project tag values in the organization's management account and create an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added. A tag policy is a type of policy that can help standardize tags across resources in the organization's accounts. A tag policy can specify the allowed tag keys, values, and case treatment for compliance. A service control policy (SCP) is a type of policy that can restrict the actions that users and roles can perform in the organization's accounts. An SCP can deny access to specific API operations unless certain conditions are met, such as having a specific tag. By creating a tag policy in the management account and attaching it to each OU, the organization can enforce consistent tagging across all accounts. By creating an SCP that denies the cloudformation:CreateStack API operation unless a project

tag is added, the organization can prevent users from creating new resources without proper tagging. This solution will meet the requirements with the least effort, as it does not involve creating additional resources or modifying existing ones.Reference:Tag policies - AWS Organizations,Service control policies - AWS Organizations,AWS CloudFormation User Guide

Question 6

Question Type: MultipleChoice

An online retail company is migrating its legacy on-premises .NET application to AWS. The application runs on load-balanced frontend web servers, load-balanced application servers, and a Microsoft SQL Server database.

The company wants to use AWS managed services where possible and does not want to rewrite the application. A solutions architect needs to implement a solution to resolve scaling issues and minimize licensing costs as the application scales.

Which solution will meet these requirements MOST cost-effectively?

Options:

A- Deploy Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer for the web tier and for the application tier. Use Amazon Aurora PostgreSQL with Babelfish turned on to replatform the SOL Server database.

B- Create images of all the servers by using AWS Database Migration Service (AWS DMS). Deploy Amazon EC2 instances that are based on the on-premises imports. Deploy the instances in an Auto Scaling group behind a Network Load Balancer for the web tier and for the application tier. Use Amazon DynamoDB as the database tier.

C- Containerize the web frontend tier and the application tier. Provision an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Create an Auto Scaling group behind a Network Load Balancer for the web tier and for the application tier. Use Amazon RDS for SOL Server to host the database.

D- Separate the application functions into AWS Lambda functions. Use Amazon API Gateway for the web frontend tier and the application tier. Migrate the data to Amazon S3. Use Amazon Athena to query the data.

Answer:

A

Explanation:

The best solution is to create a tag policy that contains the allowed project tag values in the organization's management account and create an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added. A tag policy is a type of policy that can help standardize tags across resources in the organization's accounts. A tag policy can specify the allowed tag keys, values, and case treatment for compliance. A service control policy (SCP) is a type of policy that can restrict the actions that users and roles can perform in the organization's accounts. An SCP can deny access to specific API operations unless certain conditions are met, such as having a specific tag. By creating a tag policy in the management account and attaching it to each OU, the organization can enforce consistent tagging across all accounts. By creating an SCP that denies the cloudformation:CreateStack API operation unless a project tag is added, the organization can prevent users from creating new resources without proper tagging. This solution will meet the requirements with the least effort, as it does not involve creating additional resources or modifying existing ones.Reference:Tag policies -

Question 7

Question Type: MultipleChoice

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX). and must accommodate high levels of throughput.

Which storage solution will meet these requirements?

Options:

A- Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucket. Mount the NFS file share on each EC2 instance in the duster.

B- Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mode. Mount the EFS file system on each EC2 instance in the cluster.

C- Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the io2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.

D- Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode. Mount the EFS file system on each EC2 instance in the cluster.

Answer:

D

Explanation:

The best solution is to provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode and mount the EFS file system on each EC2 instance in the cluster. Amazon EFS is a fully managed, scalable, and elastic file storage service that supports the POSIX standard and can be accessed by multiple EC2 instances concurrently. Amazon EFS offers two performance modes: General Purpose and Max I/O. Max I/O mode is designed for highly parallelized workloads that can tolerate higher latencies than the General Purpose mode. Max I/O mode provides higher levels of aggregate throughput and operations per second, which are suitable for big data analytics applications. This solution meets all the requirements of the company.Reference:Amazon EFS Documentation,Amazon EFS performance modes