



Free Questions for SAA-C03 by actualtestdumps

Shared by Whitehead on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company has a web application that is based on Java and PHP. The company plans to move the application from on-premises to AWS. The company needs the ability to test new site features frequently. The company also needs a highly available and managed solution that requires minimum operational overhead.

Which solution will meet these requirements?

Options:

- A) Create an Amazon S3 bucket. Enable static web hosting on the S3 bucket. Upload the static content to the S3 bucket. Use AWS Lambda to process all dynamic content.
- B) Deploy the web application to an AWS Elastic Beanstalk environment. Use URL swapping to switch between multiple Elastic Beanstalk environments for feature testing.
- C) Deploy the web application to Amazon EC2 instances that are configured with Java and PHP. Use Auto Scaling groups and an Application Load Balancer to manage the website's availability.
- D) Containerize the web application. Deploy the web application to Amazon EC2 instances. Use the AWS Load Balancer Controller to dynamically route traffic between containers that contain the new site features for testing.

Answer:

B

Question 2

Question Type: MultipleChoice

A developer has an application that uses an AWS Lambda function to upload files to Amazon S3 and needs the required permissions to perform the task. The developer already has an IAM user with valid IAM credentials required for Amazon S3.

What should a solutions architect do to grant the permissions?

Options:

- A) Add required IAM permissions in the resource policy of the Lambda function
- B) Create a signed request using the existing IAM credentials in the Lambda function
- C) Create a new IAM user and use the existing IAM credentials in the Lambda function.
- D) Create an IAM execution role with the required permissions and attach the IAM role to the Lambda function

Answer:

D

Question 3

Question Type: MultipleChoice

A company is building a new dynamic ordering website. The company wants to minimize server maintenance and patching. The website must be highly available and must scale read and write capacity as quickly as possible to meet changes in user demand.

Which solution will meet these requirements?

Options:

- A)** Host static content in Amazon S3 Host dynamic content by using Amazon API Gateway and AWS Lambda Use Amazon DynamoDB with on-demand capacity for the database Configure Amazon CloudFront to deliver the website content
- B)** Host static content in Amazon S3 Host dynamic content by using Amazon API Gateway and AWS Lambda Use Amazon Aurora with Aurora Auto Scaling for the database Configure Amazon CloudFront to deliver the website content
- C)** Host all the website content on Amazon EC2 instances Create an Auto Scaling group to scale the EC2 instances Use an Application Load Balancer to distribute traffic Use Amazon DynamoDB with provisioned write capacity for the database
- D)** Host all the website content on Amazon EC2 instances Create an Auto Scaling group to scale the EC2 instances Use an Application Load Balancer to distribute traffic Use Amazon Aurora with Aurora Auto Scaling for the database

Answer:

A

Question 4

Question Type: MultipleChoice

A company runs a containerised application on a Kubernetes cluster in an on premises data center. The company is using a MongoDB distribution. The company wants to migrate some of these environments to AWS but no code changes or deployment method changes are possible at this time. The company needs a solution that minimizes operational overhead.

Which solution meets these requirements?

Options:

- A)** Use Amazon Elastic Container Service (Amazon ECS) with Amazon EC2 worker nodes for compute and MongoDB on EC2 for data storage
- B)** Use Amazon Elastic Container Service (Amazon ECS) with AWS Fargate for compute and Amazon DynamoDB for data storage
- C)** Use Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 worker nodes for compute and Amazon DynamoDB for data storage
- D)** Use Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate for compute and Amazon DocumentDB (with MongoDB compatibility) for data storage

Answer:

D

Explanation:


Amazon DocumentDB (with MongoDB compatibility) is a fast, reliable, and fully managed database service. Amazon DocumentDB makes it easy to set up, operate, and scale MongoDB-compatible databases in the cloud. With Amazon DocumentDB, you can run the same application code and use the same drivers and tools that you use with MongoDB.

<https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html>

Question 5

Question Type: MultipleChoice

The following IAM policy is attached to an IAM group. This is the only policy applied to the group.



A Group members are permitted any Amazon EC2 action within the us-east-1 Region. Statements after the Allow permission are not applied.

Options:

- B)** Group members are denied any Amazon EC2 permissions in the us-east-1 Region unless they are logged in with multi-factor authentication (MFA).
- C)** Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for all Regions when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action.
- D)** Group members are allowed the ec2:StopInstances and ec2:TerminateInstances permissions for the us-east-1 Region only when logged in with multi-factor authentication (MFA). Group members are permitted any other Amazon EC2 action within the us-east-1 Region.

Answer:

D

Explanation:

This answer is correct because it reflects the effect of the IAM policy on the group members. The policy has two statements: one with an Allow effect and one with a Deny effect. The Allow statement grants permission to perform any EC2 action on any resource within the us-

east-1 Region. The Deny statement overrides the Allow statement and denies permission to perform the ec2:StopInstances and ec2:TerminateInstances actions on any resource within the us-east-1 Region, unless the group member is logged in with MF

A) Therefore, the group members can perform any EC2 action except stopping or terminating instances in the us-east-1 Region, unless they use MFA.

Question 6

Question Type: MultipleChoice

A company runs an application on a large fleet of Amazon EC2 instances. The application reads and writes entries into an Amazon DynamoDB table. The size of the DynamoDB table continuously grows but the application needs only data from the last 30 days. The company needs a solution that minimizes cost and development effort.

Which solution meets these requirements?

Options:

A) Use an AWS CloudFormation template to deploy the complete solution. Redeploy the CloudFormation stack every 30 days and delete the original stack.

- B)** Use an EC2 Instance that runs a monitoring application from AWS Marketplace. Configure the monitoring application to use Amazon DynamoDB Streams to store the timestamp when a new item is created in the table. Use a script that runs on the EC2 instance to delete items that have a timestamp that is older than 30 days.
- C)** Configure Amazon DynamoDB Streams to invoke an AWS Lambda function when a new item is created in the table. Configure the Lambda function to delete items in the table that are older than 30 days.
- D)** Extend the application to add an attribute that has a value of the current timestamp plus 30 days to each new item that is created in the table. Configure DynamoDB to use the attribute as the TTL attribute.

Answer:

D

Explanation:

Amazon DynamoDB Time to Live (TTL) allows you to define a per-item timestamp to determine when an item is no longer needed. Shortly after the date and time of the specified timestamp, DynamoDB deletes the item from your table without consuming any write throughput. TTL is provided at no extra cost as a means to reduce stored data volumes by retaining only the items that remain current for your workload's needs.

TTL is useful if you store items that lose relevance after a specific time. The following are example TTL use cases:

Remove user or sensor data after one year of inactivity in an application.

Archive expired items to an Amazon S3 data lake via Amazon DynamoDB Streams and AWS Lambda.

Retain sensitive data for a certain amount of time according to contractual or regulatory obligations.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html>

Question 7

Question Type: MultipleChoice

A company has a web application that is based on Java and PHP. The company plans to move the application from on-premises to AWS. The company needs the ability to test new site features frequently. The company also needs a highly available and managed solution that requires minimum operational overhead.

Which solution will meet these requirements?

Options:

- A)** Create an Amazon S3 bucket. Enable static web hosting on the S3 bucket. Upload the static content to the S3 bucket. Use AWS Lambda to process all dynamic content.
- B)** Deploy the web application to an AWS Elastic Beanstalk environment. Use URL swapping to switch between multiple Elastic Beanstalk environments for feature testing.
- C)** Deploy the web application to Amazon EC2 instances that are configured with Java and PHP. Use Auto Scaling groups and an

Application Load Balancer to manage the website's availability

D) Containerize the web application Deploy the web application to Amazon EC2 instances Use the AWS Load Balancer Controller to dynamically route traffic between containers that contain the new site features for testing

Answer:

B

Question 8

Question Type: MultipleChoice

An application runs on an Amazon EC2 instance that has an Elastic IP address in VPC A. The application requires access to a database in VPC B. Both VPCs are in the same AWS account.

Which solution will provide the required access MOST securely?

Options:

A) Create a DB instance security group that allows all traffic from the public IP address of the application server in VPC A.

B) Configure a VPC peering connection between VPC A and VPC B.

- C) Make the DB instance publicly accessible. Assign a public IP address to the DB instance.
- D) Launch an EC2 instance with an Elastic IP address into VPC B. Proxy all requests through the new EC2 instance.

Answer:

B

Explanation:

A VPC peering connection is a networking connection between two VPCs that enables users to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network. A VPC peering connection can be created between VPCs in the same or different AWS accounts and Regions¹. By configuring a VPC peering connection between VPC A and VPC B, the solution can provide the required access most securely.

A) Create a DB instance security group that allows all traffic from the public IP address of the application server in VPC A. This solution will not provide the required access most securely, as it involves exposing the DB instance to the public internet and relying on a single IP address for access control².

C) Make the DB instance publicly accessible. Assign a public IP address to the DB instance. This solution will not provide the required access most securely, as it involves exposing the DB instance to the public internet and allowing any source to connect to it².

D) Launch an EC2 instance with an Elastic IP address into VPC B. Proxy all requests through the new EC2 instance. This solution will not provide the required access most securely, as it involves creating an additional resource and configuring a proxy server that may

introduce latency and complexity3.

Reference URL: <https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html>

To Get Premium Files for SAA-C03 Visit

<https://www.p2pexams.com/products/saa-c03>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/saa-c03>

