SAP-C02 Q&As

AWS Certified Solutions Architect - Professional

Pass Amazon SAP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/sap-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A company is using Amazon API Gateway to deploy a private REST API that will provide access to sensitive data. The API must be accessible only from an application that is deployed in a VPC. The company deploys the API successfully. However, the API is not accessible from an Amazon EC2 instance that is deployed in the VPC.

Which solution will provide connectivity between the EC2 instance and the API?

A. Create an interface VPC endpoint for API Gateway. Attach an endpoint policy that allows apigateway:* actions. Disable private DNS naming for the VPC endpoint. Configure an API resource policy that allows access from the VPC. Use the VPC endpoint\\'s DNS name to access the API.

B. Create an interface VPC endpoint for API Gateway. Attach an endpoint policy that allows the execute- api:Invoke action. Enable private DNS naming for the VPC endpoint. Configure an API resource policy that allows access from the VPC endpoint. Use the API endpoint\\'s DNS names to access the API. Most Voted

C. Create a Network Load Balancer (NLB) and a VPC link. Configure private integration between API Gateway and the NLB. Use the API endpoint\\'s DNS names to access the API.

D. Create an Application Load Balancer (ALB) and a VPC Link. Configure private integration between API Gateway and the ALB. Use the ALB endpoint\\'s DNS name to access the API.

Correct Answer: B

According to the AWS documentation1, to access a private API from a VPC, you need to do the following: Create an interface VPC endpoint for API Gateway in your VPC. This creates a private connection between your VPC and API Gateway. Attach an endpoint policy to the VPC endpoint that allows the execute-api:Invoke action for your private API. This grants permission to invoke your API from the VPC. Enable private DNS naming for the VPC endpoint. This allows you to use the same DNS names for your private APIs as you would for public APIs. Configure a resource policy for your private API that allows access from the VPC endpoint. This controls who can access your API and under what conditions. Use the API endpoint\\'s DNS names to access the API from your VPC. For example, https://api-id.executeapi.region.amazonaws.com/stage.

QUESTION 2

A company is using Amazon OpenSearch Service to analyze data. The company loads data into an OpenSearch Service cluster with 10 data nodes from an Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1 month for read-only analysis. After 1 month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data.

The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution.

Which solution will meet these requirements MOST cost-effectively?

A. Replace all the data nodes with UltraWarm nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

B. Reduce the number of data nodes in the cluster to 2 Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Transition the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy.

C. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Add cold storage nodes to

the cluster Transition the indexes from UltraWarm to cold storage. Delete the input data from the S3 bucket after 1 month by using an S3 Lifecycle policy.

D. Reduce the number of data nodes in the cluster to 2. Add instance-backed data nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

Correct Answer: B

By reducing the number of data nodes in the cluster to 2 and adding UltraWarm nodes to handle the expected capacity, the company can reduce the cost of running the cluster. Additionally, configuring the indexes to transition to UltraWarm when OpenSearch Service ingests the data will ensure that the data is stored in the most cost- effective manner. Finally, transitioning the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy will ensure that the data is retained for compliance purposes, while also reducing the ongoing costs.

QUESTION 3

A company is migrating an application to the AWS Cloud. The application runs in an on- premises data center and writes thousands of images into a mounted NFS file system each night After the company migrates the application, the company will host the application on an Amazon EC2 instance with a mounted Amazon Elastic File System (Amazon EFS) file system.

The company has established an AWS Direct Connect connection to AWS Before the migration cutover. a solutions architect must build a process that will replicate the newly created on-premises images to the EFS file system

What is the MOST operationally efficient way to replicate the images?

A. Configure a periodic process to run the aws s3 sync command from the on-premises file system to Amazon S3 Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system

B. Deploy an AWS Storage Gateway file gateway with an NFS mount point. Mount the file gateway file system on the on-premises server. Configure a process to periodically copy the images to the mount point

C. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an S3 bucket by using public VIF. Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system.

D. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an AWS PrivateLink interface VPC endpoint for Amazon EFS by using a private VIF. Configure a DataSync scheduled task to send the images to the EFS file system every 24 hours.

Correct Answer: D

This option uses AWS DataSync to replicate the on-premises images to the EFS file system over the Direct Connect connection. AWS DataSync is a service that automates and accelerates data transfer between on-premises storage systems and AWS storage services. It can transfer data to and from Amazon EFS, Amazon FSx for Windows File Server, and Amazon S3. To use AWS DataSync, the company needs to deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. The agent connects to the AWS DataSync service endpoint in the AWS Region where the EFS file system is located. The company can use an AWS PrivateLink interface endpoint to connect to the service endpoint securely and privately over the Direct Connect connection. The company can then create a task in AWS DataSync that specifies the source location (the NFS file system), the destination location (the EFS file system), and the options for the data transfer (such as schedule, bandwidth limit, and verification). AWS DataSync will then perform the data transfer efficiently and securely, using encryption in transit and at rest.

QUESTION 4

A retail company wants to improve its application architecture. The company\\'s applications register new orders, handle returns of merchandise, and provide analytics. The applications store retail data in a MySQL database and an Oracle OLAP analytics database. All the applications and databases are hosted on Amazon EC2 instances.

Each application consists of several components that handle different parts of the order process. These components use incoming data from different sources. A separate ETL job runs every week and copies data from each application to the analytics database.

A solutions architect must redesign the architecture into an event-driven solution that uses serverless services. The solution must provide updated analytics in near real time.

Which solution will meet these requirements?

A. Migrate the individual applications as microservices to Amazon Elastic Container Service (Amazon ECS) containers that use AWS Fargate. Keep the retail MySQL database on Amazon EC2. Move the analytics database to Amazon Neptune. Use Amazon Simple Queue Service (Amazon SQS) to send all the incoming data to the microservices and the analytics database.

B. Create an Auto Scaling group for each application. Specify the necessary number of EC2 instances in each Auto Scaling group. Migrate the retail MySQL database and the analytics database to Amazon Aurora MySQL. Use Amazon Simple Notification Service (Amazon SNS) to send all the incoming data to the correct EC2 instances and the analytics database.

C. Migrate the individual applications as microservices to Amazon Elastic Kubernetes Service (Amazon EKS) containers that use AWS Fargate. Migrate the retail MySQL database to Amazon Aurora Serverless MySQL. Migrate the analytics database to Amazon Redshift Serverless. Use Amazon EventBridge to send all the incoming data to the microservices and the analytics database.

D. Migrate the individual applications as microservices to Amazon AppStream 2.0. Migrate the retail MySQL database to Amazon Aurora MySQL. Migrate the analytics database to Amazon Redshift Serverless. Use AWS IoT Core to send all the incoming data to the microservices and the analytics database.

Correct Answer: C

QUESTION 5

A solutions architect has implemented a SAML 2.0 federated identity solution with their company\\'s on-premises identity provider (IdP) to authenticate users\\' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal access to the AWS environment is granted However, when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment.

Which items should the solutions architect check to ensure identity federation is property configured? (Select THREE j

A. The IAM user\\'s permissions pokey has allowed the use of SAML federation for that user

B. The IAM roles created for the federated users\\' or federated groups\\' trust policy have set the SAML provider as the principle.

C. Test users are not in the AWSFederatedUsers group in the company\\'s IdP

D. The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider the ARN of the

IAM role, and the SAML assertion from IdP

E. The on-premises IdP\\'s DNS hostname is reachable from the AWS environment VPCs.

F. The company\\'s IdP defines SAML assertions that property map users or groups m the company to IAM roles with appropriate permissions

Correct Answer: BCE