

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.leads4pass.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 has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an Application Load Balancer (ALB) that is associated with three public subnets.

The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.

A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically.

Which solution will meet these requirements?

- A. Deploy three NAT gateways, one in each public subnet. Assign the Elastic IP address to the NAT gateways. Turn on health checks for the NAT gateways. If a NAT gateway fails a health check, recreate the NAT gateway and assign the Elastic IP address to the new NAT gateway.
- B. Replace the ALB with a Network Load Balancer (NLB). Assign the Elastic IP address to the NLB. Turn on health checks for the NLB. In the case of a failed health check, redeploy the NLB in different subnets.
- C. Deploy a single NAT gateway in a public subnet. Assign the Elastic IP address to the NAT gateway. Use Amazon CloudWatch with a custom metric to monitor the NAT gateway. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subnet. Assign the Elastic IP address to the new NAT gateway.
- D. Assign the Elastic IP address to the ALB. Create an Amazon Route 53 simple record with the Elastic IP address as the value. Create a Route 53 health check. In the case of a failed health check, recreate the ALB in different subnets.

Correct Answer: C

to connect out from the private subnet you need an NAT gateway and since only one Elastic IP whitelisted on firewall its one NATGateway at time and if AZ failure happens Lambda creates a new NATGATEWAY in a different AZ using the Same Elastic IP ,dont be tempted to select D since application that needs to connect is on a private subnet whose outbound connections use the NATGateway Elastic IP

QUESTION 2

A retail company runs a business-critical web service on an Amazon Elastic Container Service (Amazon ECS) cluster that runs on Amazon EC2 instances. The web service receives POST requests from end users and writes data to a MySQL database that runs on a separate EC2 instance. The company needs to ensure that data loss does not occur. The current code deployment process includes manual updates of the ECS service. During a recent deployment, end users encountered intermittent 502 Bad Gateway errors in response to valid web requests.

The company wants to implement a reliable solution to prevent this issue from recurring. The company also wants to automate code deployments. The solution must be highly available and must optimize cost-effectiveness.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS

CodeDeploy to perform a blue/green deployment with validation testing to update the ECS service.

B. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses Provisioned IOPS SSD (io2) storage

C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source to receive the POST requests from the web service Configure an AWS Lambda function to poll the queue Write the data to the database.

D. Run the web service on an ECS cluster that has a Fargate launch type Use AWS CodePipeline and AWS CodeDeploy to perform a canary deployment to update the ECS service.

Correct Answer: CD

QUESTION 3

A company manages hundreds of AWS accounts centrally in an organization In AWS Organizations. The company recently started to allow product teams to create and manage their own S3 access points in their accounts. The S3 access points can be accessed only within VPCs. not on the internet.

What is the MOST operationally efficient way to enforce this requirement?

A. Set the S3 access point resource policy to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.

B. Create an SCP at the root level in the organization to deny the s3: Create Access Point action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.

C. Use AWS Cloud Formation StackSets to create a new IAM policy In each AWS account that allows the s3:CreateAccessPoint action only if the s3:AccessPointNetworkOrigin condition key evaluates to VPC.

D. Set the S3 bucket policy to deny the s3:CreateAccessPoint action unless the s3: AccessPointNetworkOrigin condition key evaluates to VPC.

Correct Answer: B

<https://aws.amazon.com/s3/features/access-points/> <https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/>

QUESTION 4

A company is planning to set up a REST API application on AWS. The application team wants to set up a new identity store on AWS The IT team does not want to maintain any infrastructure or servers for this deployment.

What is the MOST operationally efficient solution that meets these requirements?

A. Deploy the application as AWS Lambda functions. Set up Amazon API Gateway REST API endpoints for the application Create a Lambda function, and configure a Lambda authorizer

B. Deploy the application in AWS AppSync, and configure AWS Lambda resolvers Set up an Amazon Cognito user pool, and configure AWS AppSync to use the user pool for authorization C. Deploy the application as AWS Lambda functions. Set up Amazon API Gateway REST API endpoints for the application Set up an Amazon Cognito user pool, and configure an Amazon Cognito authorizer

D. Deploy the application in Amazon Elastic Kubernetes Service (Amazon EKS) clusters. Set up an Application Load Balancer for the EKS pods Set up an Amazon Cognito user pool and service pod for authentication.

Correct Answer: C

QUESTION 5

A company deploys a new web application As part of the setup, the company configures AWS WAF to log to Amazon S3 through Amazon Kinesis Data Firehose. The company develops an Amazon Athena query that runs once daily to return AWS WAF log data from the previous 24 hours. The volume of daily logs is constant However over time, the same query is taking more time to run

A solutions architect needs to design a solution to prevent the query time from continuing to increase. The solution must minimize operational overhead

Which solution will meet these requirements?

- A. Create an AWS Lambda function that consolidates each day's AWS WAF logs into one log file
- B. Reduce the amount of data scanned by configuring AWS WAF to send logs to a different S3 bucket each day
- C. Update the Kinesis Data Firehose configuration to partition the data in Amazon S3 by date and time Create external tables for Amazon Redshift Configure Amazon Redshift Spectrum to query the data source
- D. Modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and time. Change the Athena query to view the relevant partitions

Correct Answer: D

The best solution is to modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and time. This will reduce the amount of data scanned by Athena and improve the query performance. Changing

the Athena query to view the relevant partitions will also help to filter out unnecessary data. This solution requires minimal operational overhead as it does not involve creating additional resources or changing the log format.

References:

[AWS WAF Developer Guide], [Amazon Kinesis Data Firehose User Guide], [Amazon Athena User Guide]

[SAP-C02 VCE Dumps](#)

[SAP-C02 Study Guide](#)

[SAP-C02 Braindumps](#)