ANS-C01$^{Q\&As}$

AWS Certified Advanced Networking Specialty Exam

Pass Amazon ANS-C01 Exam with 100% Guarantee

QUESTION 1

A company has an application that runs on a fleet of Amazon EC2 instances. A new company regulation mandates that all network traffic toand from the EC2 instances must be sent to a centralized third-party EC2 appliance for content inspection.Which solution will meet these requirements?

A. Configure VPC flow logs on each EC2 network interface. Publish the flow logs to an Amazon S3 bucket. Create a third-party EC2appliance to acquire flow logs from the S3 bucket. Log in to the appliance to monitor network content.

B. Create a third-party EC2 appliance in an Auto Scaling group fronted by a Network Load Balancer (NLB). Configure a mirror session.Specify the NLB as the mirror target. Specify a mirror filter to capture inbound and outbound traffic. For the source of the mirror session,specify the EC2 elastic network interfaces for all the instances that host the application.

C. Configure a mirror session. Specify an Amazon Kinesis Data Firehose delivery stream as the mirror target. Specify a mirror filter tocapture inbound and outbound traffic. For the source of the mirror session, specify the EC2 elastic network interfaces for all the instancesthat host the application. Create a third-party EC2 appliance. Send all traffic to the appliance through the Kinesis Data Firehose deliverystream for content inspection.

D. Configure VPC flow logs on each EC2 network interface. Send the logs to Amazon CloudWatch. Create a third-party EC2 appliance.Configure a CloudWatch filter to send the flow logs to Amazon Kinesis Data Firehose to load the logs into the appliance.

Correct Answer: B

You can use the following resources as traffic mirror targets: Network interfaces of type interface Network Load Balancers Gateway Load Balancer endpoints https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-targets.html

QUESTION 2

A company is developing an application in which IoT devices will report measurements to the AWS Cloud. The application will have millions ofend users. The company observes that the IoT devices cannot support DNS resolution. The company needs to implement an Amazon EC2 AutoScaling solution so that the IoT devices can connect to an application endpoint without using DNS.Which solution will meet these requirements MOST cost-effectively?

A. Use an Application Load Balancer (ALB)-type target group for a Network Load Balancer (NLB). Create an EC2 Auto Scaling group. Attachthe Auto Scaling group to the ALB. Set up the IoT devices to connect to the IP addresses of the NLB.

B. Use an AWS Global Accelerator accelerator with an Application Load Balancer (ALB) endpoint. Create an EC2 Auto Scaling group. Attachthe Auto Scaling group to the ALSet up the IoT devices to connect to the IP addresses of the accelerator.

C. Use a Network Load Balancer (NLB). Create an EC2 Auto Scaling group. Attach the Auto Scaling group to the NLB. Set up the IoTdevices to connect to the IP addresses of the NLB.

D. Use an AWS Global Accelerator accelerator with a Network Load Balancer (NLB) endpoint. Create an EC2 Auto Scaling group. Attach theAuto Scaling group to the NLB. Set up the IoT devices to connect to the IP addresses of the accelerator.

Correct Answer: C

B, C, and D are also doable.

Let\\'s think about the cost.

AWS Global Accelerator is definitely the best option. but it costs more money.

NLB is enough.

QUESTION 3

A company is planning to create a service that requires encryption in transit. The traffic must not be decrypted between the client and thebackend of the service. The company will implement the service by using the gRPC protocol over TCP port 443. The service will scale up tothousands of simultaneous connections. The backend of the service will be hosted on an Amazon Elastic Kubernetes Service (Amazon EKS)duster with the Kubernetes Cluster Autoscaler and the Horizontal Pod Autoscaler configured. The company needs to use mutual TLS for two-way authentication between the client and the backend.Which solution will meet these requirements?

A. Install the AWS Load Balancer Controller for Kubernetes. Using that controller, configure a Network Load Balancer with a TCP listeneron port 443 to forward traffic to the IP addresses of the backend service Pods.

B. Install the AWS Load Balancer Controller for Kubernetes. Using that controller, configure an Application Load Balancer with an HTTPSlistener on port 443 to forward traffic to the IP addresses of the backend service Pods.

C. Create a target group. Add the EKS managed node group\\'s Auto Scaling group as a target Create an Application Load Balancer with anHTTPS listener on port 443 to forward traffic to the target group.

D. Create a target group. Add the EKS managed node group\\'s Auto Scaling group as a target. Create a Network Load Balancer with a TLSlistener on port 443 to forward traffic to the target group.

Correct Answer: A

ALB does support HTTP/2 and gRPC workloads. However, the title mentions that the company needs to use mutual TLS for mutual authentication between the client and the backend. This means that traffic cannot be decrypted between the client and the service backend. Since the ALB will terminate the TLS connection and decrypt the traffic, it does not meet the requirements in the title. In contrast, NLB can forward TCP traffic without decrypting the traffic, so it is more suitable for meeting the needs described in the title.

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/configure-mutual-tls-authentication-for-applications-running-on-amazon-eks.html

QUESTION 4

A network engineer is working on a large migration effort from an on-premises data center to an AWS Control Tower based multi-account environment. The environment has a transit gateway that is deployed to a central network services

account. The central network services account has been shared with an organization in AWS Organizations through AWS Resource Access Manager (AWS RAM).

A shared services account also exists in the environment. The shared services account hosts workloads that need to be shared with the entire organization.

The network engineer needs to create a solution to automate the deployment of common network components across the environment. The solution must provision a VPC for application workloads to each new and existing member account.

The VPCs must be connected to the transit gateway in the central network services account.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose three.)

A. Deploy an AWS Lambda function to the shared services account. Program the Lambda function to assume a role in the new and existing member accounts to provision the necessary network infrastructure.

B. Update the existing accounts with an Account Factory Customization (AFC). Select the same AFC when provisioning new accounts.

C. Create an AWS CloudFormation template that describes the infrastructure that needs to be created in each account. Upload the template as an AWS Service Catalog product to the shared services account.

D. Deploy an Amazon EventBridge rule on a default event bus in the shared services account. Configure the EventBridge rule to react to AWS Control Tower CreateManagedAccount lifecycle events and to invoke the AWS Lambda function.

E. Create an AWSControlTowerBiueprintAccess role in the shared services account. F Create an AWSControlTowerBiueprintAccess role in each member account.

Correct Answer: BCE

QUESTION 5

A company has business operations in the United States and in Europe. The company\\'s public applications are running on AWS and use threetransit gateways. The transit gateways are located in the us-west-2, us-east-1, and eu-central-1 Regions. All the transit gateways areconnected to each other in a full mesh configuration.The company accidentally removes the route to the eu-central-1 VPCs from the us-west-2 transit gateway route table. The company alsoaccidentally removes the route to the us-west-2 VPCs from the eu-central-1 transit gateway route table.How can a network engineer identify the misconfiguration with the LEAST operational overhead?

A. Use the Route Analyzer feature for AWS Transit Gateway Network Manager.

B. Use the AWSSupport-SetupIPMonitoringFromVPC AWS Systems Manager Automation runbook. Push network telemetry data to AmazonCloudWatch Logs for analysis.

C. Use VPC flow logs in eu-central-1 and us-west-2 to analyze the missing routes.

D. Use Amazon VPC Traffic Mirroring in eu-central-1 or us-west-2 to take packet captures and troubleshoot the connectivity issues.

Correct Answer: A

https://docs.aws.amazon.com/network-manager/latest/tgwnm/route-analyzer.html

Latest ANS-C01 Dumps ANS-C01 VCE Dumps ANS-C01 Exam Questions