



# ANS-C01<sup>Q&As</sup>

AWS Certified Advanced Networking Specialty Exam

## Pass Amazon ANS-C01 Exam with 100% Guarantee

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

<https://www.passapply.com/ans-c01.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 developed an application on AWS that will track inventory levels of vending machines and initiate the restocking process automatically. The company plans to integrate this application with vending machines and deploy the vending machines in several markets around the world. The application resides in a VPC in the us-east-1 Region. The application consists of an Amazon Elastic Container Service (Amazon ECS) cluster behind an Application Load Balancer (ALB). The communication from the vending machines to the application happens over HTTPS. The company is planning to use an AWS Global Accelerator accelerator and configure static IP addresses of the accelerator in the vending machines for application endpoint access. The application must be accessible only through the accelerator and not through a direct connection over the internet to the ALB endpoint. Which solution will meet these requirements?

- A. Configure the ALB in a private subnet of the VPC. Attach an internet gateway without adding routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- B. Configure the ALB in a private subnet of the VPC. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- C. Configure the ALB in a public subnet of the VPC. Attach an internet gateway. Add routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.
- D. Configure the ALB in a private subnet of the VPC. Attach an internet gateway. Add routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.

Correct Answer: A

This is not a normal scenario of attaching IGW to EC2 instance by creating a route in subnet.

Please read the below link typically describing ELB integration with AWS Global accelerator (and the last line of the extract)

<https://docs.aws.amazon.com/global-accelerator/latest/dg/secure-vpc-connections.html>

"When you add an internal Application Load Balancer or an Amazon EC2 instance endpoint in AWS Global Accelerator, you enable internet traffic to flow directly to and from the endpoint in Virtual Private Clouds (VPCs) by targeting it in a private subnet. The VPC that contains the load balancer or EC2 instance must have an internet gateway attached to it, to indicate that the VPC accepts internet traffic. However, you don't need public IP addresses on the load balancer or EC2 instance. You also don't need an associated internet gateway route for the subnet."

## QUESTION 2

An application team for a startup company is deploying a new multi-tier application into the AWS Cloud. The application will be hosted on a fleet of Amazon EC2 instances that run in an Auto Scaling group behind a publicly accessible Network Load Balancer (NLB). The application requires the clients to work with UDP traffic and TCP traffic. In the near term, the application will serve only users within the same geographic location. The application team plans to extend the application to a global audience and will move the deployment to multiple AWS Regions around the world to bring the application closer to the end users. The application team wants to use the new Regions to deploy new versions of the application and wants to be able to control the amount of traffic that each Region receives during these rollouts. In addition, the application team must minimize first-byte latency and jitter (randomized delay) for the end users. How should the application team design the network architecture for the application to meet these requirements?



- A. Create an Amazon CloudFront distribution to align to each Regional deployment. Set the NLB for each Region as the origin for each CloudFront distribution. Use an Amazon Route 53 weighted routing policy to control traffic to the newer Regional deployments.
- B. Create an AWS Global Accelerator accelerator and listeners for the required ports. Configure endpoint groups for each Region. Configure a traffic dial for the endpoint groups to control traffic to the newer Regional deployments. Register the NLBs with the endpoint groups.
- C. Use Amazon S3 Transfer Acceleration for the application in each Region. Adjust the amount of traffic that each Region receives from the Transfer Acceleration endpoints to the Regional NLBs.
- D. Create an Amazon CloudFront distribution that includes an origin group. Set the NLB for each Region as the origins for the origin group. Use an Amazon Route 53 latency routing policy to control traffic to the new Regional deployments.

Correct Answer: B

CloudFront is designed to handle HTTP protocol meanwhile Global Accelerator is best used for both HTTP and non-HTTP protocols such as TCP and UDP. and CloudFront doesn't support NLB

### QUESTION 3

A network engineer needs to build an encrypted connection between an on-premises data center and a VPC. The network engineer attaches the VPC to a virtual private gateway and sets up an AWS Site-to-Site VPN connection. The VPN tunnel is UP after configuration and is working. However, during rekey for phase 2 of the VPN negotiation, the customer gateway device is receiving different parameters than the parameters that the device is configured to support. The network engineer checks the IPsec configuration of the VPN tunnel. The network engineer notices that the customer gateway device is configured with the most secure encryption algorithms that the AWS Site-to-Site VPN configuration file provides. What should the network engineer do to troubleshoot and correct the issue?

- A. Check the native virtual private gateway logs. Restrict the VPN tunnel options to the specific VPN parameters that the virtual private gateway requires.
- B. Check the native customer gateway logs. Restrict the VPN tunnel options to the specific VPN parameters that the customer gateway requires.
- C. Check Amazon CloudWatch logs of the virtual private gateway. Restrict the VPN tunnel options to the specific VPN parameters that the virtual private gateway requires.
- D. Check Amazon CloudWatch logs of the customer gateway. Restrict the VPN tunnel options to the specific VPN parameters that the customer gateway requires.

Correct Answer: B

You check Cloudwatch for AWS resources or your native/on-prem logs for your on-prem resource. A and D is out.

The problem statement indicates that customer gateway is misconfigured. So you need to work on Customer gateway.

### QUESTION 4

A company is deploying third-party firewall appliances for traffic inspection and NAT capabilities in its VPC. The VPC is configured with private subnets and public subnets. The company needs to deploy the firewall appliances behind a load balancer. Which architecture will meet these requirements MOST cost-effectively?



- A. Deploy a Gateway Load Balancer with the firewall appliances as targets. Configure the firewall appliances with a single network interface in a private subnet. Use a NAT gateway to send the traffic to the internet after inspection.
- B. Deploy a Gateway Load Balancer with the firewall appliances as targets. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subnet. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.
- C. Deploy a Network Load Balancer with the firewall appliances as targets. Configure the firewall appliances with a single network interface in a private subnet. Use a NAT gateway to send the traffic to the internet after inspection.
- D. Deploy a Network Load Balancer with the firewall appliances as targets. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subnet. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.

Correct Answer: B

Gateway Load balancer, use the built in NAT functionality of the firewall to save money and two network interfaces to inspect both private and public subnets.

---

#### QUESTION 5

A company's application team is unable to launch new resources into its VPC. A network engineer discovers that the VPC has run out of usable IP addresses. The VPC CIDR block is 172.16.0.0/16. Which additional CIDR block can the network engineer attach to the VPC?

- A. 172.17.0.0/29
- B. 10.0.0.0/16
- C. 172.17.0.0/16
- D. 192.168.0.0/16

Correct Answer: C

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-cidr-blocks.html#add-cidr-block-restrictions> Only it will also allow us to add 172.17.0.0/16

[ANS-C01 PDF Dumps](#)

[ANS-C01 VCE Dumps](#)

[ANS-C01 Practice Test](#)