Amazon

ANS-C01
AWS Certified Advanced Networking - Specialty

Exams Boost

Boost Up Your Career

For More Information – Visit link below:

https://www.examsboost.com/
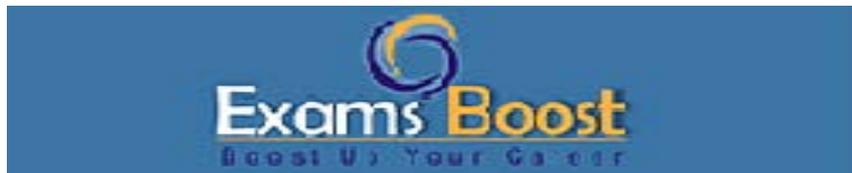
Product Version

✓ **Up to Date products, reliable and verified.**
✓ **Questions and Answers in PDF Format.**

Latest Version: 9.0

Question: 1

A company is planning to create a service that requires encryption in transit. The traffic must not be decrypted between the client and the backend of the service. The company will implement the service by using the gRPC protocol over TCP port 443. The service will scale up to thousands 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 listener on 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 HTTPS listener 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 an HTTPS 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 TLS listener on port 443 to forward traffic to the target group.

Answer: B

Explanation:
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-targetgroups. html#target-group-protocol-version https://docs.aws.amazon.com/prescriptiveguidance/ latest/patterns/deploy-a-grpc-based-application-on-an-amazon-eks-cluster-and-access-itwith-an-application-load-balancer.html

Question: 2

A company is deploying a new application in the AWS Cloud. The company wants a highly available web server that will sit behind an Elastic Load Balancer. The load balancer will route requests to multiple target groups based on the URL in the request. All traffic must use HTTPS. TLS processing must be offloaded to the load balancer. The web server must know the user's IP address so that the company can keep accurate logs for security purposes.
Which solution will meet these requirements?

A. Deploy an Application Load Balancer with an HTTPS listener. Use path-based routing rules to forward the traffic to the correct target group. Include the X-Forwarded-For request header with traffic to the

targets.

B. Deploy an Application Load Balancer with an HTTPS listener for each domain. Use host-based routing rules to forward the traffic to the correct target group for each domain. Include the X-Forwarded-For request header with traffic to the targets.

C. Deploy a Network Load Balancer with a TLS listener. Use path-based routing rules to forward the traffic to the correct target group. Configure client IP address preservation for traffic to the targets.

D. Deploy a Network Load Balancer with a TLS listener for each domain. Use host-based routing rules to forward the traffic to the correct target group for each domain. Configure client IP address preservation for traffic to the targets.

Answer: A

Explanation:

An Application Load Balancer (ALB) can be used to route traffic to multiple target groups based on the URL in the request. The ALB can be configured with an HTTPS listener to ensure all traffic uses HTTPS. TLS processing can be offloaded to the ALB, which reduces the load on the web server. Path-based routing rules can be used to route traffic to the correct target group based on the URL in the request. The XForwarded-

For request header can be included with traffic to the targets, which will allow the web server to know the user's IP address and keep accurate logs for security purposes.

Question: 3

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 VPAttach 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.

Answer: A

Explanation:
Please read the below link typically describing ELB integration with AWS Global accelator (and the last line of the extract) - https://docs.aws.amazon.com/global-accelerator/latest/dg/secure-vpcconnections. 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: 4

A global delivery company is modernizing its fleet management system. The company has several business units. Each business unit designs and maintains applications that are hosted in its own AWS account in separate application VPCs in the same AWS Region. Each business unit's applications are designed to get data from a central shared services VPC.
The company wants the network connectivity architecture to provide granular security controls. The architecture also must be able to scale as more business units consume data from the central shared services VPC in the future.
Which solution will meet these requirements in the MOST secure manner?

A. Create a central transit gateway. Create a VPC attachment to each application VPC. Provide full mesh connectivity between all the VPCs by using the transit gateway.
B. Create VPC peering connections between the central shared services VPC and each application VPC in each business unit's AWS account.
C. Create VPC endpoint services powered by AWS PrivateLink in the central shared services VPCCreate VPC endpoints in each application VPC.
D. Create a central transit VPC with a VPN appliance from AWS Marketplace. Create a VPN attachment from each VPC to the transit VPC. Provide full mesh connectivity among all the VPCs.

Answer: C

Explanation:
Option C provides a secure and scalable solution using VPC endpoint services powered by AWS PrivateLink. AWS PrivateLink enables private connectivity between VPCs and services without exposing the data to the public internet or using a VPN connection. By creating VPC endpoints in each application VPC, the company can securely access the central shared services VPC without the need for complex network configurations. Furthermore, PrivateLink supports cross-account connectivity, which makes it a

scalable solution as more business units consume data from the central shared services VPC in the future.

Question: 5

A company uses a 4 Gbps AWS Direct Connect dedicated connection with a link aggregation group (LAG) bundle to connect to five VPCs that are deployed in the us-east-1 Region. Each VPC serves a different business unit and uses its own private VIF for connectivity to the on-premises environment. Users are reporting slowness when they access resources that are hosted on AWS.
A network engineer finds that there are sudden increases in throughput and that the Direct Connect connection becomes saturated at the same time for about an hour each business day. The company wants to know which business unit is causing the sudden increase in throughput. The network engineer must find out this information and implement a solution to resolve the problem.
Which solution will meet these requirements?

A. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observed. Create a new 10 Gbps dedicated connection. Shift traffic from the existing dedicated connection to the new dedicated connection.
B. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observed. Upgrade the bandwidth of the existing dedicated connection to 10 Gbps.
C. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observed. Upgrade the existing dedicated connection to a 5 Gbps hosted connection.
D. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observed. Create a new 10 Gbps dedicated connection. Shift traffic from the existing dedicated connection to the new dedicated connection.

Answer: A

Explanation:
To meet the requirements of finding out which business unit is causing the sudden increase in throughput and resolving the problem, the network engineer should review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observed (Option B). After identifying the VIF that is causing the issue, they can upgrade the bandwidth of the existing dedicated connection to 10 Gbps to resolve the problem (Option B).

Thank You for Trying Our Product

For More Information – Visit link below:

https://www.examsboost.com/

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ **90 Days Free Updates**
- ✓ **Money Back Pass Guarantee**
- ✓ **Instant Download or Email Attachment**
- ✓ **24/7 Live Chat Support**
- ✓ **PDF file could be used at any Platform**
- ✓ **50,000 Happy Customer**