

Amazon.ANS-C01.v2023-07-13.q34

Exam Code:	ANS-C01
Exam Name:	AWS Certified Advanced Networking Specialty Exam
Certification Provider:	Amazon
Free Question Number:	34
Version:	v2023-07-13
# of views:	515
# of Questions views:	340
https://www.freepdfdumps.com/Amazon.ANS-C01.v2023-07-13.q34.html	

NEW QUESTION: 1

A company is running multiple workloads on Amazon EC2 instances in public subnets. In a recent incident, an attacker exploited an application vulnerability on one of the EC2 instances to gain access to the instance. The company fixed the application and launched a replacement EC2 instance that contains the updated application.

The attacker used the compromised application to spread malware over the internet. The company became aware of the compromise through a notification from AWS. The company needs the ability to identify when an application that is deployed on an EC2 instance is spreading malware.

Which solution will meet this requirement with the LEAST operational effort?

- A. Use Amazon GuardDuty to analyze traffic patterns by inspecting DNS requests and VPC flow logs.
- B. Use Amazon GuardDuty to deploy AWS managed decoy systems that are equipped with the most recent malware signatures.
- C. Set up a Gateway Load Balancer. Run an intrusion detection system (IDS) appliance from AWS Marketplace on Amazon EC2 for traffic inspection.
- D. Configure Amazon Inspector to perform deep packet inspection of outgoing traffic.

Answer: A (LEAVE A REPLY)

This solution involves using Amazon GuardDuty to monitor network traffic and analyze DNS requests and VPC flow logs for suspicious activity. This will allow the company to identify when an application is spreading malware by monitoring the network traffic patterns associated with the instance. GuardDuty is a fully managed threat detection service that continuously monitors for malicious activity and unauthorized behavior in your AWS accounts and workloads. It requires minimal setup and configuration and can be integrated with other AWS services for automated remediation. This solution requires the least operational effort compared to the other options

NEW QUESTION: 2

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 (LEAVE A REPLY)

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).

NEW QUESTION: 3

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 VPC. Create 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 (LEAVE A REPLY)

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.

NEW QUESTION: 4

A banking company is successfully operating its public mobile banking stack on AWS. The mobile banking stack is deployed in a VPC that includes private subnets and public subnets. The company is using IPv4 networking and has not deployed or supported IPv6 in the environment. The company has decided to adopt a third-party service provider's API and must integrate the API with the existing environment. The service provider's API requires the use of IPv6.

A network engineer must turn on IPv6 connectivity for the existing workload that is deployed in a private subnet. The company does not want to permit IPv6 traffic from the public internet and mandates that the company's servers must initiate all IPv6 connectivity. The network engineer turns on IPv6 in the VPC and in the private subnets.

Which solution will meet these requirements?

A. Create an internet gateway and a NAT instance in the VPC. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT instance.

B. Create an egress-only internet gateway in the VPC. Configure a security group that denies all inbound traffic. Associate the security group with the egress-only internet gateway.

C. Create an egress-only Internet gateway in the VPAdd a route to the existing subnet route tables to point IPv6 traffic to the egress-only internet gateway.

D. Create an internet gateway and a NAT gateway in the VPC. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT gateway.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 5

A company's network engineer is designing an active-passive connection to AWS from two on-premises data centers. The company has set up AWS Direct Connect connections between the on-premises data centers and AWS. From each location, the company is using a transit VIF that connects to a Direct Connect gateway that is associated with a transit gateway.

The network engineer must ensure that traffic from AWS to the data centers is routed first to the primary data center. The traffic should be routed to the failover data center only in the case of an outage.

Which solution will meet these requirements?

A. Set the BGP community tag for all prefixes from the primary data center to 7224:7300.

Set the BGP community tag for all prefixes from the failover data center to 7224:7100

B. Set the BGP community tag for all prefixes from the primary data center to 7224:9300.

Set the BGP community tag for all prefixes from the failover data center to 7224:9100

C. Set the BGP community tag for all prefixes from the primary data center to 7224:7100.

Set the BGP community tag for all prefixes from the failover data center to 7224:7300

D. Set the BGP community tag for all prefixes from the primary data center to 7224:9100.

Set the BGP community tag for all prefixes from the failover data center to 7224:9300

Answer: **A** ([LEAVE A REPLY](#))

NEW QUESTION: 6

A company uses AWS Direct Connect to connect its corporate network to multiple VPCs in the same AWS account and the same AWS Region. Each VPC uses its own private VIF and its own virtual LAN on the Direct Connect connection. The company has grown and will soon surpass the limit of VPCs and private VIFs for each connection.

What is the MOST scalable way to add VPCs with on-premises connectivity?

A. Provision a new Direct Connect connection to handle the additional VPCs. Use the new connection to connect additional VPCs.

B. Create virtual private gateways for each VPC that is over the service quota. Use AWS Site-to-Site VPN to connect the virtual private gateways to the corporate network.

C. Create a Direct Connect gateway, and add virtual private gateway associations to the VPCs. Configure a private VIF to connect to the corporate network.

D. Create a transit gateway, and attach the VPCs. Create a Direct Connect gateway, and associate it with the transit gateway. Create a transit VIF to the Direct Connect gateway.

Answer: ([SHOW ANSWER](#))

When a company requires connectivity to multiple VPCs over AWS Direct Connect, a scalable solution is to use a transit gateway. A transit gateway is a hub that can interconnect multiple VPCs and VPN connections. The VPCs can communicate with each other over the transit gateway, and on-premises networks can communicate with the VPCs through the Direct Connect gateway. This solution provides a central point of management and simplifies the configuration of network routing. By associating the Direct Connect gateway with the transit gateway, traffic between the VPCs and the on-premises network can be routed through the Direct Connect connection.

NEW QUESTION: 7

A company is deploying an application. The application is implemented in a series of containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use the Fargate launch type for its tasks. The containers will run workloads that require connectivity initiated over an SSL connection. Traffic must be able to flow to the application from other AWS accounts over private connectivity. The application must scale in a manageable way as more consumers use the application.

Which solution will meet these requirements?

A. Choose a Network Load Balancer (NLB) as the type of load balancer for the ECS service. Specify the NLB in the service definition. Create a VPC endpoint service for the NLB. Share the VPC endpoint service with other AWS accounts.

B. Choose a Gateway Load Balancer (GLB) as the type of load balancer for the ECS service. Create a lifecycle hook to add new tasks to the target group from Amazon ECS as required to handle scaling. Specify the GLB in the service definition. Create a VPC peer for external AWS accounts. Update the route tables so that the AWS accounts can reach the GLB.

C. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service. Create path-based routing rules to allow the application to target the containers that are registered in the target group. Specify the ALB in the service definition. Create a VPC peer for the external AWS accounts. Update the route tables so that the AWS accounts can reach the ALB.

D. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service. Create path-based routing rules to allow the application to target the containers that are registered in the target group. Specify the ALB in the service definition. Create a VPC endpoint service for the ALB. Share the VPC endpoint service with other AWS accounts.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 8

A company's network engineer is designing a hybrid DNS solution for an AWS Cloud workload. Individual teams want to manage their own DNS hostnames for their applications in their development environment. The solution must integrate the application-specific hostnames with the centrally managed DNS hostnames from the on-premises network and must provide bidirectional name resolution. The solution also must minimize management overhead.

Which combination of steps should the network engineer take to meet these requirements? (Choose three.)

- A.** Create DNS proxy servers.
- B.** Use an Amazon Route 53 Resolver outbound endpoint.
- C.** Create Amazon Route 53 private hosted zones.
- D.** Modify the DHCP options set by setting a custom DNS server value.
- E.** Use an Amazon Route 53 Resolver inbound endpoint.
- F.** Set up a zone transfer between Amazon Route 53 and the on-premises DNS.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 9

A media company is implementing a news website for a global audience. The website uses Amazon CloudFront as its content delivery network. The backend runs on Amazon EC2 Windows instances behind an Application Load Balancer (ALB). The instances are part of an Auto Scaling group. The company's customers access the website by using service.example.com as the CloudFront custom domain name. The CloudFront origin points to an ALB that uses service-alb.example.com as the domain name.

The company's security policy requires the traffic to be encrypted in transit at all times between the users and the backend.

Which combination of changes must the company make to meet this security requirement? (Choose three.)

- A.** Create a certificate for service.example.com by using AWS Certificate Manager (ACM). Configure CloudFront to use this custom SSL/TLS certificate. Change the default behavior to redirect HTTP to HTTPS.
- B.** Create a public certificate from a third-party certificate provider with any domain name for the EC2 instances. Configure the backend to use this certificate for its HTTPS listener. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its targets. Attach the existing Auto Scaling group to this new target group.
- C.** Create a self-signed certificate for service-alb.example.com. Import the certificate into AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the imported service-alb.example.com ACM certificate. Modify the CloudFront origin to use the HTTPS protocol only. Delete the HTTP listener on the ALB.
- D.** Create a certificate for service-alb.example.com by using AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the

service-alb.example.com ACM certificate. Modify the CloudFront origin to use the HTTPS protocol only. Delete the HTTP listener on the ALB.

E. Create a self-signed certificate for service.example.com. Import the certificate into AWS Certificate Manager (ACM). Configure CloudFront to use this imported SSL/TLS certificate. Change the default behavior to redirect HTTP to HTTPS.

F. Create a certificate with any domain name by using AWS Certificate Manager (ACM) for the EC2 instances. Configure the backend to use this certificate for its HTTPS listener. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its targets. Attach the existing Auto Scaling group to this new target group.

Answer: A,B,D ([LEAVE A REPLY](#))

NEW QUESTION: 10

A network engineer must provide additional safeguards to protect encrypted data at Application Load Balancers (ALBs) through the use of a unique random session key. What should the network engineer do to meet this requirement?

- A.** Change the ALB security policy to a policy that supports TLS 1.2 protocol only
- B.** Use AWS Key Management Service (AWS KMS) to encrypt session keys
- C.** Associate an AWS WAF web ACL with the ALBs. and create a security rule to enforce forward secrecy (FS)
- D.** Change the ALB security policy to a policy that supports forward secrecy (FS)

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 11

A company has two on-premises data center locations. There is a company-managed router at each data center. Each data center has a dedicated AWS Direct Connect connection to a Direct Connect gateway through a private virtual interface. The router for the first location is advertising 110 routes to the Direct Connect gateway by using BGP, and the router for the second location is advertising 60 routes to the Direct Connect gateway by using BGP. The Direct Connect gateway is attached to a company VPC through a virtual private gateway.

A network engineer receives reports that resources in the VPC are not reachable from various locations in either data center. The network engineer checks the VPC route table and sees that the routes from the first data center location are not being populated into the route table. The network engineer must resolve this issue in the most operationally efficient manner.

What should the network engineer do to meet these requirements?

- A.** Remove the Direct Connect gateway, and create a new private virtual interface from each company router to the virtual private gateway of the VPC.
- B.** Change the router configurations to summarize the advertised routes.
- C.** Open a support ticket to increase the quota on advertised routes to the VPC route table.

D. Create an AWS Transit Gateway. Attach the transit gateway to the VPC, and connect the Direct Connect gateway to the transit gateway.

Answer: B (LEAVE A REPLY)

"If you advertise more than 100 routes each for IPv4 and IPv6 over the BGP session, the BGP session will go into an idle state with the BGP session DOWN."

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html>

NEW QUESTION: 12

A security team is performing an audit of a company's AWS deployment. The security team is concerned that two applications might be accessing resources that should be blocked by network ACLs and security groups. The applications are deployed across two Amazon Elastic Kubernetes Service (Amazon EKS) clusters that use the Amazon VPC Container Network Interface (CNI) plugin for Kubernetes. The clusters are in separate subnets within the same VPC and have a Cluster Autoscaler configured.

The security team needs to determine which POD IP addresses are communicating with which services throughout the VPC. The security team wants to limit the number of flow logs and wants to examine the traffic from only the two applications.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create VPC flow logs in a custom format. Set the application subnets as resources. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.

B. Create VPC flow logs in a custom format. Create a filter to gather flow logs only from the EKS nodes. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.

C. Create VPC flow logs in the default format. Create a filter to gather flow logs only from the EKS nodes. Include the srcaddr field and the dstaddr field in the flow logs.

D. Create VPC flow logs in a custom format. Set the EKS nodes as the resource. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 13

A company has deployed Amazon EC2 instances in private subnets in a VPC. The EC2 instances must initiate any requests that leave the VPC, including requests to the company's on-premises data center over an AWS Direct Connect connection. No resources outside the VPC can be allowed to open communications directly to the EC2 instances.

The on-premises data center's customer gateway is configured with a stateful firewall device that filters for incoming and outgoing requests to and from multiple VPCs. In addition, the company wants to use a single IP match rule to allow all the communications from the EC2 instances to its data center from a single IP address.

Which solution will meet these requirements with the LEAST amount of operational overhead?

- A.** Configure the on-premises firewall to filter all requests from the on-premises network to the EC2 instances. Allow a stateful connection if the EC2 instances in the VPC initiate the traffic.
- B.** Deploy a NAT instance into a private subnet in the VPC where the EC2 instances are deployed. Configure the on-premises firewall to allow connections from the IP address that is assigned to the NAT instance.
- C.** Deploy a NAT gateway into a private subnet in the VPC where the EC2 instances are deployed. Specify the NAT gateway type as private. Configure the on-premises firewall to allow connections from the IP address that is assigned to the NAT gateway.
- D.** Create a VPN connection over the Direct Connect connection by using the on-premises firewall. Use the firewall to block all traffic from on premises to AWS. Allow a stateful connection from the EC2 instances to initiate the requests.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 14

A company is planning to deploy many software-defined WAN (SD-WAN) sites. The company is using AWS Transit Gateway and has deployed a transit gateway in the required AWS Region. A network engineer needs to deploy the SD-WAN hub virtual appliance into a VPC that is connected to the transit gateway. The solution must support at least 5 Gbps of throughput from the SD-WAN hub virtual appliance to other VPCs that are attached to the transit gateway.

Which solution will meet these requirements?

- A.** Assign a new CIDR block to the transit gateway. Create a new VPC for the SD-WAN hub virtual appliance. Attach the new VPC to the transit gateway with a VPC attachment. Add a transit gateway Connect attachment. Create a Connect peer and specify the VXLAN and BGP parameters. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.
- B.** Assign a new CIDR block to the transit gateway. Create a new VPC for the SD-WAN hub virtual appliance. Attach the new VPC to the transit gateway with a VPC attachment. Add a transit gateway Connect attachment. Create a Connect peer and specify the GRE and BGP parameters. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.
- C.** Create a new VPC for the SD-WAN hub virtual appliance. Attach the new VPC to the transit gateway with a VPC attachment. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway. Configure BGP over the IPsec VPN connections.
- D.** Create a new VPC for the SD-WAN hub virtual appliance. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway. Configure BGP over the IPsec VPN connections

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 15

A company has deployed an AWS Network Firewall firewall into a VPC. A network engineer needs to implement a solution to deliver Network Firewall flow logs to the company's Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster in the shortest possible time.

Which solution will meet these requirements?

- A.** Create an Amazon S3 bucket. Create an AWS Lambda function to load logs into the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster. Enable Amazon Simple Notification Service (Amazon SNS) notifications on the S3 bucket to invoke the Lambda function. Configure flow logs for the firewall. Set the S3 bucket as the destination.
- B.** Create an Amazon Kinesis Data Firehose delivery stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination. Configure flow logs for the firewall. Set the Kinesis Data Firehose delivery stream as the destination for the Network Firewall flow logs.
- C.** Configure flow logs for the firewall. Set the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination for the Network Firewall flow logs.
- D.** Create an Amazon Kinesis data stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination. Configure flow logs for the firewall. Set the Kinesis data stream as the destination for the Network Firewall flow logs.

Answer: B (LEAVE A REPLY)

<https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-analyze-aws-network-firewall-logs-using-amazon-opensearch-service-part-1/>

NEW QUESTION: 16

An Australian ecommerce company hosts all of its services in the AWS Cloud and wants to expand its customer base to the United States (US). The company is targeting the western US for the expansion.

The company's existing AWS architecture consists of four AWS accounts with multiple VPCs deployed in the ap-southeast-2 Region. All VPCs are attached to a transit gateway in ap-southeast-2. There are dedicated VPCs for each application service. The company also has VPCs for centralized security features such as proxies, firewalls, and logging. The company plans to duplicate the infrastructure from ap-southeast-2 to the us-west-1 Region. A network engineer must establish connectivity between the various applications in the two Regions. The solution must maximize bandwidth, minimize latency and minimize operational overhead.

Which solution will meet these requirements?

- A.** Create VPN attachments between the two transit gateways. Configure the VPN attachments to use BGP routing between the two transit gateways.
- B.** Peer the transit gateways in each Region. Configure routing between the two transit gateways for each Region's IP addresses.

C. Create a VPN server in a VPC in each Region. Update the routing to point to the VPN servers for the IP addresses in alternate Regions.

D. Attach the VPCs in us-west-1 to the transit gateway in ap-southeast-2.

Answer: B (LEAVE A REPLY)

Peering the transit gateways in each region would establish a private network connection between the two regions, allowing the company to route traffic between the VPCs in different regions without going over the public internet. This would help minimize latency and maximize bandwidth while reducing the operational overhead of managing multiple VPN connections.

Valid ANS-C01 Dumps shared by Actual4test.com for Helping Passing ANS-C01 Exam! Actual4test.com now offer the **newest ANS-C01 exam dumps**, the Actual4test.com ANS-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com ANS-C01 dumps with Test Engine here: https://www.actual4test.com/ANS-C01_examcollection.html (291 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 17

A company has two AWS accounts one for Production and one for Connectivity. A network engineer needs to connect the Production account VPC to a transit gateway in the Connectivity account. The feature to auto accept shared attachments is not enabled on the transit gateway.

Which set of steps should the network engineer follow in each AWS account to meet these requirements?

A. 1. In the Production account: Create a resource share in AWS Resource Access Manager for the transit gateway. Provide the Connectivity account ID. Enable the feature to allow external accounts

2. In the Connectivity account: Accept the resource.

3. In the Connectivity account: Create an attachment to the VPC subnets.

4. In the Production account: Accept the attachment. Associate a route table with the attachment.

B. 1. In the Production account: Create a resource share in AWS Resource Access Manager for the VPC subnets. Provide the Connectivity account ID. Enable the feature to allow external accounts.

2. In the Connectivity account: Accept the resource.

3. In the Production account: Create an attachment on the transit gateway to the VPC subnets.

4. In the Connectivity account: Accept the attachment. Associate a route table with the attachment.

C. 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the VPC subnets. Provide the Production account ID. Enable the feature to allow external accounts.

2. In the Production account: Accept the resource.

3. In the Connectivity account: Create an attachment on the transit gateway to the VPC subnets.

4. In the Production account: Accept the attachment. Associate a route table with the attachment.

D. 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the transit gateway. Provide the Production account ID. Enable the feature to allow external accounts.

2. In the Production account: Accept the resource.

3. In the Production account: Create an attachment to the VPC subnets.

4. In the Connectivity account: Accept the attachment. Associate a route table with the attachment.

Answer: A (LEAVE A REPLY)

step 1: In the Production account, create a resource share in AWS Resource Access Manager for the transit gateway and provide the Connectivity account ID. Enabling the feature to allow external accounts is also required to share resources between accounts.

Step 2: In the Connectivity account, accept the shared resource. This action will allow the Production account to use the transit gateway in the Connectivity account. Step 3: In the Connectivity account, create an attachment to the VPC subnets. This attachment will enable communication between the VPC in the Production account and the transit gateway in the Connectivity account. Step 4: In the Production account, accept the attachment and associate a route table with the attachment. This will enable the VPC to route traffic through the transit gateway to other resources in the Connectivity account.

NEW QUESTION: 18

A company is using custom DNS servers that run BIND for name resolution in its VPCs. The VPCs are deployed across multiple AWS accounts that are part of the same organization in AWS Organizations. All the VPCs are connected to a transit gateway. The BIND servers are running in a central VPC and are configured to forward all queries for an on-premises DNS domain to DNS servers that are hosted in an on-premises data center. To ensure that all the VPCs use the custom DNS servers, a network engineer has configured a VPC DHCP options set in all the VPCs that specifies the custom DNS servers to be used as domain name servers.

Multiple development teams in the company want to use Amazon Elastic File System (Amazon EFS). A development team has created a new EFS file system but cannot mount the file system to one of its Amazon EC2 instances. The network engineer discovers that the EC2 instance cannot resolve the IP address for the EFS mount point `fs-33444567d.efs.us-east-1.amazonaws.com`. The network engineer needs to implement a

solution so that development teams throughout the organization can mount EFS file systems.

Which combination of steps will meet these requirements? (Choose two.)

- A.** Configure the BIND DNS servers in the central VPC to forward queries for `efs.us-east-1.amazonaws.com` to the Amazon provided DNS server (169.254.169.253).
- B.** Create an Amazon Route 53 Resolver outbound endpoint in the central VPC. Update all the VPC DHCP options sets to use AmazonProvidedDNS for name resolution.
- C.** Create an Amazon Route 53 Resolver inbound endpoint in the central VPC. Update all the VPC DHCP options sets to use the Route 53 Resolver inbound endpoint in the central VPC for name resolution.
- D.** Create an Amazon Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS servers. Share the rule with the organization by using AWS Resource Access Manager (AWS RAM). Associate the rule with all the VPCs.
- E.** Create an Amazon Route 53 private hosted zone for the `efs.us-east-1.amazonaws.com` domain. Associate the private hosted zone with the VPC where the EC2 instance is deployed. Create an A record for `fs-33444567d.efs.us-east-1.amazonaws.com` in the private hosted zone. Configure the A record to return the mount target of the EFS mount point.

Answer: ([SHOW ANSWER](#))

Option B suggests using Amazon Route 53 Resolver outbound endpoint, which would replace the existing BIND DNS servers with the AmazonProvidedDNS for name resolution. However, the scenario specifically mentions that the company is using custom DNS servers that run BIND for name resolution in its VPCs, so this solution would not work. Option D suggests creating a Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS servers, which would not address the issue of resolving the EFS mount point. The problem is not with resolving queries for the on-premises domain, but rather with resolving the IP address for the EFS mount point.

NEW QUESTION: 19

Your security team implements a host-based firewall on all of your Amazon Elastic Compute Cloud (EC2) instances to block all outgoing traffic. Exceptions must be requested for each specific requirement. Until you request a new rule, you cannot access the instance metadata service. Which firewall rule should you request to be added to your instances to allow instance metadata access?

- A.** Inbound; Protocol tcp; Source [Instance's EIP]; Destination 169.254.169.254
- B.** Inbound; Protocol tcp; Destination 169.254.169.254; Destination port 80
- C.** Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 80
- D.** Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 443

Answer: **C** ([LEAVE A REPLY](#))

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html>
To view all categories of instance metadata from within a running instance, use the following URI. <http://169.254.169.254/latest/meta-data/>

NEW QUESTION: 20

A Network Engineer is provisioning a subnet for a load balancer that will sit in front of a fleet of application servers in a private subnet. There is limited IP space left in the VPC CIDR. The application has few users now but is expected to grow quickly to millions of users.

What design will use the LEAST amount of IP space, while allowing for this growth?

- A.** Use two /29 subnets for an Application Load Balancer in different Availability Zones.
- B.** Use one /29 subnet for the Network Load Balancer. Add another VPC CIDR to the VPC to allow for future growth.
- C.** Use one /28 subnet for an Application Load Balancer. Add another VPC CIDR to the VPC to allow for future growth.
- D.** Use two /28 subnets for a Network Load Balancer in different Availability Zones.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 21

A company's development team has created a new product recommendation web service. The web service is hosted in a VPC with a CIDR block of 192.168.224.0/19. The company has deployed the web service on Amazon EC2 instances and has configured an Auto Scaling group as the target of a Network Load Balancer (NLB).

The company wants to perform testing to determine whether users who receive product recommendations spend more money than users who do not receive product recommendations. The company has a big sales event in 5 days and needs to integrate its existing production environment with the recommendation engine by then. The existing production environment is hosted in a VPC with a CIDR block of 192.168.128.0/17.

A network engineer must integrate the systems by designing a solution that results in the least possible disruption to the existing environments.

Which solution will meet these requirements?

- A.** Ask the development team of the web service to redeploy the web service into the production VPC and integrate the systems there.
- B.** Create a transit gateway in the existing production environment. Create attachments to the production VPC and the web service VPC. Configure appropriate routing rules in the transit gateway and VPC route tables for 192.168.224.0/19 and 192.168.128.0/17. Configure the relevant security groups and ACLs to allow the systems to communicate.
- C.** Create a VPC peering connection between the web service VPC and the existing production VPC. Add a routing rule to the appropriate route table to allow data to flow to 192.168.224.0/19 from the existing production environment and to flow to

192.168.128.0/17 from the web service environment. Configure the relevant security groups and ACLs to allow the systems to communicate.

D. Create a VPC endpoint service. Associate the VPC endpoint service with the NLB for the web service. Create an interface VPC endpoint for the web service in the existing production VPC.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 22

A network engineer needs to update a company's hybrid network to support IPv6 for the upcoming release of a new application. The application is hosted in a VPC in the AWS Cloud. The company's current AWS infrastructure includes VPCs that are connected by a transit gateway. The transit gateway is connected to the on-premises network by AWS Direct Connect and AWS Site-to-Site VPN. The company's on-premises devices have been updated to support the new IPv6 requirements.

The company has enabled IPv6 for the existing VPC by assigning a new IPv6 CIDR block to the VPC and by assigning IPv6 to the subnets for dual-stack support. The company has launched new Amazon EC2 instances for the new application in the updated subnets.

When updating the hybrid network to support IPv6 the network engineer must avoid making any changes to the current infrastructure. The network engineer also must block direct access to the instances' new IPv6 addresses from the internet. However, the network engineer must allow outbound internet access from the instances.

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

A. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Create a new VPN connection that supports IPv6 connectivity. Add a NAT gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.

B. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Update the existing VPN connection to support IPv6 connectivity. Add an egress-only internet gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.

C. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Create a new VPN connection that supports IPv6 connectivity. Add an egress-only internet gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices

D. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Create a new VPN connection that supports IPv6 connectivity. Add an egress-only internet gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 23

A network engineer needs to set up an Amazon EC2 Auto Scaling group to run a Linux-based network appliance in a highly available architecture. The network engineer is configuring the new launch template for the Auto Scaling group.

In addition to the primary network interface the network appliance requires a second network interface that will be used exclusively by the application to exchange traffic with hosts over the internet. The company has set up a Bring Your Own IP (BYOIP) pool that includes an Elastic IP address that should be used as the public IP address for the second network interface.

How can the network engineer implement the required architecture?

A. Configure the two network interfaces in the launch template. Define the primary network interface to be created in one of the private subnets. For the second network interface, select one of the public subnets. Choose the BYOIP pool ID as the source of public IP addresses.

B. Configure the primary network interface in a private subnet in the launch template. Use the user data option to run a cloud-init script after boot to attach the second network interface from a subnet with auto-assign public IP addressing enabled.

C. Create an AWS Lambda function to run as a lifecycle hook of the Auto Scaling group when an instance is launching. In the Lambda function, assign a network interface to an AWS Global Accelerator endpoint.

D. During creation of the Auto Scaling group, select subnets for the primary network interface. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

Answer: ([SHOW ANSWER](#))

During creation of the Auto Scaling group, select subnets for the primary network interface. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

This solution meets all of the requirements stated in the question. The primary network interface can be configured in a private subnet during creation of the Auto Scaling group. The user data option can be used to run a cloud-init script that will allocate a second network interface and associate an Elastic IP address from the BYOIP pool with it.

NEW QUESTION: 24

A company is hosting an application on Amazon EC2 instances behind a Network Load Balancer (NLB). A solutions architect added EC2 instances in a second Availability Zone to improve the availability of the application. The solutions architect added the instances to the NLB target group.

The company's operations team notices that traffic is being routed only to the instances in the first Availability Zone.

What is the MOST operationally efficient solution to resolve this issue?

- A. Enable the new Availability Zone on the NLB
- B. Create a new NLB for the instances in the second Availability Zone
- C. Enable proxy protocol on the NLB
- D. Create a new target group with the instances in both Availability Zones

Answer: A (LEAVE A REPLY)

When adding instances in a new Availability Zone to an existing Network Load Balancer (NLB), it is important to ensure that the new Availability Zone is enabled on the NLB. This will allow traffic to be routed to instances in both Availability Zones. This can be done by editing the settings of the NLB and selecting the new Availability Zone from the list of available zones.

NEW QUESTION: 25

A company plans to deploy a two-tier web application to a new VPC in a single AWS Region. The company has configured the VPC with an internet gateway and four subnets. Two of the subnets are public and have default routes that point to the internet gateway. Two of the subnets are private and share a route table that does not have a default route. The application will run on a set of Amazon EC2 instances that will be deployed behind an external Application Load Balancer. The EC2 instances must not be directly accessible from the internet. The application will use an Amazon S3 bucket in the same Region to store data. The application will invoke S3 GET API operations and S3 PUT API operations from the EC2 instances. A network engineer must design a VPC architecture that minimizes data transfer cost.

Which solution will meet these requirements?

- A. Deploy the EC2 instances in the public subnets. Create an S3 interface endpoint in the VPC. Modify the application configuration to use the S3 endpoint-specific DNS hostname.
- B. Deploy the EC2 instances in the private subnets. Create a NAT gateway in the VPC. Create default routes in the private subnets to the NAT gateway. Connect to Amazon S3 by using the NAT gateway.
- C. Deploy the EC2 instances in the private subnets. Create an S3 gateway endpoint in the VPC. Specify the route table of the private subnets during endpoint creation to create routes to Amazon S3.
- D. Deploy the EC2 instances in the private subnets. Create an S3 interface endpoint in the VPC. Modify the application configuration to use the S3 endpoint-specific DNS hostname.

Answer: C (LEAVE A REPLY)

Option C is the optimal solution as it involves deploying the EC2 instances in the private subnets, which provides additional security benefits. Additionally, creating an S3 gateway endpoint in the VPC will enable the EC2 instances to communicate with Amazon S3 directly, without incurring data transfer costs. This is because the S3 gateway endpoint uses Amazon's private network to transfer data between the VPC and S3, which is not charged for data transfer. Furthermore, specifying the route table of the private subnets

during endpoint creation will create routes to Amazon S3, which is required for the EC2 instances to communicate with S3.

NEW QUESTION: 26

A company uses a hybrid architecture and has an AWS Direct Connect connection between its on-premises data center and AWS. The company has production applications that run in the on-premises data center. The company also has production applications that run in a VPC. The applications that run in the on-premises data center need to communicate with the applications that run in the VPC. The company is using corp.example.com as the domain name for the on-premises resources and is using an Amazon Route 53 private hosted zone for aws.example.com to host the VPC resources. The company is using an open-source recursive DNS resolver in a VPC subnet and is using a DNS resolver in the on-premises data center. The company's on-premises DNS resolver has a forwarder that directs requests for the aws.example.com domain name to the DNS resolver in the VPC. The DNS resolver in the VPC has a forwarder that directs requests for the corp.example.com domain name to the DNS resolver in the on-premises data center. The company has decided to replace the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints.

Which combination of steps should a network engineer take to make this replacement? (Choose three.)

- A. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the outbound endpoint.
- B. Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.
- C. Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint.
- D. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.
- E. Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver.
- F. Configure the on-premises DNS resolver to forward aws.example.com queries to the IP addresses of the outbound endpoint.

Answer: B,C,E (LEAVE A REPLY)

To replace the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints in a hybrid architecture where on-premises applications need to communicate with applications running in a VPC, a network engineer should take the following steps: Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint. (Option C) Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint. (Option B) Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver. (Option E) These steps will allow for seamless

replacement of the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints and enable communication between on-premises and VPC applications.

NEW QUESTION: 27

You deploy an Amazon EC2 instance that runs a web server into a subnet in a VPC. An Internet gateway is attached, and the main route table has a default route (0.0.0.0/0) configured with a target of the Internet gateway.

The instance has a security group configured to allow as follows:

Protocol: TCP

Port: 80 inbound, nothing outbound

The Network ACL for the subnet is configured to allow as follows:

Protocol: TCP

Port: 80 inbound, nothing outbound

When you try to browse to the web server, you receive no response.

Which additional step should you take to receive a successful response?

- A. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 80
- B. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 1024-65535
- C. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 80
- D. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 1024-65535

Answer: D (LEAVE A REPLY)

To enable the connection to a service running on an instance, the associated network ACL must allow both inbound traffic on the port that the service is listening on as well as allow outbound traffic from ephemeral ports. When a client connects to a server, a random port from the ephemeral port range (1024-65535) becomes the client's source port. The designated ephemeral port then becomes the destination port for return traffic from the service, so outbound traffic from the ephemeral port must be allowed in the network ACL.

<https://aws.amazon.com/premiumsupport/knowledge-center/resolve-connection-sg-acl-inbound/>

NEW QUESTION: 28

A real estate company is building an internal application so that real estate agents can upload photos and videos of various properties. The application will store these photos and videos in an Amazon S3 bucket as objects and will use Amazon DynamoDB to store corresponding metadata. The S3 bucket will be configured to publish all PUT events for new object uploads to an Amazon Simple Queue Service (Amazon SQS) queue.

A compute cluster of Amazon EC2 instances will poll the SQS queue to find out about newly uploaded objects. The cluster will retrieve new objects, perform proprietary image and video recognition and classification, update metadata in DynamoDB, and replace the

objects with new watermarked objects. The company does not want public IP addresses on the EC2 instances.

Which networking design solution will meet these requirements MOST cost-effectively as application usage increases?

A. Place the EC2 instances in a private subnet. Create a NAT gateway in a public subnet in the same Availability Zone. Create an internet gateway. Attach the internet gateway to the VPC. In the public subnet's route table, add a default route that points to the internet gateway

B. Place the EC2 instances in a private subnet. Create a gateway VPC endpoint for Amazon SQS. Create interface VPC endpoints for Amazon S3 and DynamoDB.

C. Place the EC2 instances in a private subnet. Create an interface VPC endpoint for Amazon SQS. Create gateway VPC endpoints for Amazon S3 and DynamoDB.

D. Place the EC2 instances in a public subnet. Disable the Auto-assign Public IP option while launching the EC2 instances. Create an internet gateway. Attach the internet gateway to the VPC. In the public subnet's route table, add a default route that points to the internet gateway.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 29

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.

Answer: A (LEAVE A REPLY)

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."

NEW QUESTION: 30

A company deploys a new web application on Amazon EC2 instances. The application runs in private subnets in three Availability Zones behind an Application Load Balancer (ALB). Security auditors require encryption of all connections. The company uses Amazon Route 53 for DNS and uses AWS Certificate Manager (ACM) to automate SSL/TLS certificate provisioning. SSL/TLS connections are terminated on the ALB.

The company tests the application with a single EC2 instance and does not observe any problems. However, after production deployment, users report that they can log in but that they cannot use the application. Every new web request restarts the login process.

What should a network engineer do to resolve this issue?

A. Modify the ALB listener configuration. Edit the rule that forwards traffic to the target group. Change the rule to enable group-level stickiness. Set the duration to the maximum application session length.

B. Modify the ALB target group configuration by enabling the stickiness attribute. Use an application-based cookie. Set the duration to the maximum application session length.

C. Remove the ALB. Create an Amazon Route 53 rule with a failover routing policy for the application name. Configure ACM to issue certificates for each EC2 instance.

D. Replace the ALB with a Network Load Balancer. Create a TLS listener. Create a new target group with the protocol type set to TLS Register the EC2 instances. Modify the target group configuration by enabling the stickiness attribute.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 31

Your company runs an application for the US market in the us-east-1 AWS region. This application uses proprietary TCP and UDP protocols on Amazon Elastic Compute Cloud (EC2) instances. End users run a real-time, front-end application on their local PCs. This front-end application knows the DNS hostname of the service.

You must prepare the system for global expansion. The end users must access the application with lowest latency.

How should you use AWS services to meet these requirements?

- A.** Register the IP addresses of the service hosts as "A" records with latency-based routing policy in Amazon Route 53, and set a Route 53 health check for these hosts.
- B.** Set Amazon CloudFront in front of the host of the service, and register the CloudFront name of the main service as an ALIAS record in Route 53.
- C.** Set the Amazon API gateway in front of the service, and register the API gateway name of the main service as an ALIAS record in Route 53.
- D.** Set the Elastic Load Balancing (ELB) load balancer in front of the hosts of the service, and register the ELB name of the main service host as an ALIAS record with a latency-based routing policy in Route 53.

Answer: D (LEAVE A REPLY)

Valid ANS-C01 Dumps shared by Actual4test.com for Helping Passing ANS-C01 Exam! Actual4test.com now offer the **newest ANS-C01 exam dumps**, the Actual4test.com ANS-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com ANS-C01 dumps with Test Engine here: https://www.actual4test.com/ANS-C01_examcollection.html (291 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 32

A network engineer is designing the architecture for a healthcare company's workload that is moving to the AWS Cloud. All data to and from the on-premises environment must be encrypted in transit. All traffic also must be inspected in the cloud before the traffic is allowed to leave the cloud and travel to the on-premises environment or to the internet. The company will expose components of the workload to the internet so that patients can reserve appointments. The architecture must secure these components and protect them against DDoS attacks. The architecture also must provide protection against financial liability for services that scale out during a DDoS event.

Which combination of steps should the network engineer take to meet all these requirements for the workload? (Choose three.)

- A.** Use Traffic Mirroring to copy all traffic to a fleet of traffic capture appliances.
- B.** Set up AWS WAF on all network components.

- C.** Configure an AWS Lambda function to create Deny rules in security groups to block malicious IP addresses.
- D.** Use AWS Direct Connect with MACsec support for connectivity to the cloud.
- E.** Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection.
- F.** Configure AWS Shield Advanced and ensure that it is configured on all public assets.

Answer: ([SHOW ANSWER](#))

To meet the requirements for the healthcare company's workload that is moving to the AWS Cloud, the network engineer should take the following steps:

Use AWS Direct Connect with MACsec support for connectivity to the cloud to ensure that all data to and from the on-premises environment is encrypted in transit (Option D).

Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection to inspect all traffic in the cloud before it is allowed to leave (Option E).

Configure AWS Shield Advanced and ensure that it is configured on all public assets to secure components exposed to the internet against DDoS attacks and provide protection against financial liability for services that scale out during a DDoS event (Option F).

These steps will help ensure that all data is encrypted in transit, all traffic is inspected before leaving the cloud, and components exposed to the internet are secured against DDoS attacks.

NEW QUESTION: 33

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 ([LEAVE A REPLY](#))

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 X-Forwarded-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.

NEW QUESTION: 34

All IP addresses within a 10.0.0.0/16 VPC are fully utilized with application servers across two Availability Zones. The application servers need to send frequent UDP probes to a single central authentication server on the Internet to confirm that is running up-to-date packages. The network is designed for application servers to use a single NAT gateway for internal access. Testing reveals that a few of the servers are unable to communicate with the authentication server.

- A.** The NAT gateway does not support UDP traffic.
- B.** The authentication server is not accepting traffic.
- C.** The NAT gateway cannot allocate more ports.
- D.** The NAT gateway is launched in a private subnet.

Answer: C (LEAVE A REPLY)

Ref: <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

"A NAT gateway can support up to 55,000 simultaneous connections to each unique destination. This limit also applies if you create approximately 900 connections per second to a single destination (about 55,000 connections per minute). If the destination IP address, the destination port, or the protocol (TCP/UDP/ICMP) changes, you can create an additional 55,000 connections. For more than 55,000 connections, there is an increased chance of connection errors due to port allocation errors. These errors can be monitored by viewing the ErrorPortAllocation CloudWatch metric for your NAT gateway. For more information, see Monitoring NAT Gateways Using Amazon CloudWatch."

Valid ANS-C01 Dumps shared by Actual4test.com for Helping Passing ANS-C01 Exam! Actual4test.com now offer the **newest ANS-C01 exam dumps**, the Actual4test.com ANS-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com ANS-C01 dumps with Test Engine here: https://www.actual4test.com/ANS-C01_examcollection.html (291 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)