Amazon-Web-Services

Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam

NEW QUESTION 1
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 forefs.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 VP
C. Update all the VPC DHCP options sets to use AmazonProvidedDNS for name resolution.
D. Create an Amazon Route 53 Resolver inbound endpoint in the central VPUpdate all the VPC DHCP options sets to use the Route 53 Resolver inbound endpoint in the central VPC for name resolution.
E. Create an Amazon Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS server
F. Share the rule with the organization by using AWS Resource Access Manager (AWS RAM). Associate the rule with all the VPCs.
G. 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 deploye
H. Create an A record for fs-33444567d.efs.us-east-1.amazonaws.com in the private hosted zon
I. Configure the A record to return the mount target of the EFS mount point.

Answer: BD

Explanation:
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 2
A network engineer needs to standardize a company's approach to centralizing and managing interface VPC endpoints for private communication with AWS services. The company uses AWS Transit Gateway for inter-VPC connectivity between AWS accounts through a hub-and-spoke model. The company's network services team must manage all Amazon Route 53 zones and interface endpoints within a shared services AWS account. The company wants to use thiscentralized model to provide AWS resources with access to AWS Key Management Service (AWS KMS) without sending traffic over the public internet.
What should the network engineer do to meet these requirements?

A. In the shared services account, create an interface endpoint for AWS KM
B. Modify the interface endpoint by disabling the private DNS nam
C. Create a private hosted zone in the shared services account with an alias record that points to the interface endpoin
D. Associate the private hosted zone with the spoke VPCs in each AWS account.
E. In the shared services account, create an interface endpoint for AWS KM
F. Modify the interface endpoint by disabling the private DNS nam
G. Create a private hosted zone in each spoke AWS account with an alias record that points to the interface endpoin
H. Associate each private hosted zone with the shared services AWS account.
I. In each spoke AWS account, create an interface endpoint for AWS KM
J. Modify each interface endpoint by disabling the private DNS nam
K. Create a private hosted zone in each spoke AWS account with an alias record that points to each interface endpoin
L. Associate each private hosted zone with the shared services AWS account.
M. In each spoke AWS account, create an interface endpoint for AWS KM
N. Modify each interface endpoint by disabling the private DNS nam
O. Create a private hosted zone in the shared services account with an alias record that points to each interface endpoin
P. Associate the private hosted zone with the spoke VPCs in each AWS account.

Answer: A

NEW QUESTION 3
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. Use an Amazon Route 53 Resolver inbound endpoint.
B. Modify the DHCP options set by setting a custom DNS server value.
C. Use an Amazon Route 53 Resolver outbound endpoint.
D. Create DNS proxy servers.
E. Create Amazon Route 53 private hosted zones.
F. Set up a zone transfer between Amazon Route 53 and the on-premises DNS.

Answer: ABE

NEW QUESTION 4
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 templat
B. Define the primary network interface to be created in one of the private subnet
C. For the second network interface, select one of the public subnet
D. Choose the BYOIP pool ID as the source of public IP addresses.
E. Configure the primary network interface in a private subnet in the launch templat
F. 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.
G. Create an AWS Lambda function to run as a lifecycle hook of the Auto Scaling group when an instance is launchin
H. In the Lambda function, assign a network interface to an AWS Global Accelerator endpoint.
I. During creation of the Auto Scaling group, select subnets for the primary network interfac
J. Use the user data option to run a cloud-init script to allocate a second network interface and to associate anElastic IP address from the BYOIP pool.

Answer: D

Explanation:
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 5
A software-as-a-service (SaaS) provider hosts its solution on Amazon EC2 instances within a VPC in the AWS Cloud. All of the provider's customers also have their environments in the AWS Cloud.
A recent design meeting revealed that the customers have IP address overlap with the provider's AWS deployment. The customers have stated that they will not share their internal IP addresses and that they do not want to connect to the provider's SaaS service over the internet.
Which combination of steps is part of a solution that meets these requirements? (Choose two.)

A. Deploy the SaaS service endpoint behind a Network Load Balancer.
B. Configure an endpoint service, and grant the customers permission to create a connection to the endpoint service.
C. Deploy the SaaS service endpoint behind an Application Load Balancer.
D. Configure a VPC peering connection to the customer VPC
E. Route traffic through NAT gateways.
F. Deploy an AWS Transit Gateway, and connect the SaaS VPC to i
G. Share the transit gateway with the customer
H. Configure routing on the transit gateway.

Answer: AB

Explanation:
NLB for creating the private link which solves the overlapping IP address issue and the SaaS service endpoint behind it. (the SaaS endpoint could be an ALB)https://aws.amazon.com/about-aws/whats-new/2021/09/application-load-balancer-aws-privatelink-static-ip

NEW QUESTION 6
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 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.
C. 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.
D. 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.

Answer: B

NEW QUESTION 7
A company has deployed a software-defined WAN (SD-WAN) solution to interconnect all of its offices. The company is migrating workloads to AWS and needs to extend its SD-WAN solution to support connectivity to these workloads.
A network engineer plans to deploy AWS Transit Gateway Connect and two SD-WAN virtual appliances to provide this connectivity. According to company policies, only a single SD-WAN virtual appliance can handle traffic from AWS workloads at a given time.
How should the network engineer configure routing to meet these requirements?

A. Add a static default route in the transit gateway route table to point to the secondary SD-WAN virtual applianc
B. Add routes that are more specific to point to the primary SD-WAN virtual appliance.
C. Configure the BGP community tag 7224:7300 on the primary SD-WAN virtual appliance for BGP routes toward the transit gateway.
D. Configure the AS_PATH prepend attribute on the secondary SD-WAN virtual appliance for BGP routes toward the transit gateway.
E. Disable equal-cost multi-path (ECMP) routing on the transit gateway for Transit Gateway Connect.

Answer: A

NEW QUESTION 8

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 Gateway Load Balancer (GLB) as the type of load balancer for the ECS servic
B. Create a lifecycle hook to add new tasks to the target group from Amazon ECS as required to handle scalin
C. Specify the GLB in the service definitio
D. Create a VPC peer for external AWS account
E. Update the route tables so that the AWS accounts can reach the GLB.
F. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS servic
G. Create path-based routing rules to allow the application to target the containers that are registered in the target grou
H. Specify the ALB in the service definitio
I. Create a VPC endpoint service for the ALB Share the VPC endpoint service with other AWS accounts.
J. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS servic
K. Create path-based routing rules to allow the application to target the containers that are registered in the target grou
L. Specify the ALB in the service definitio
M. Create a VPC peer for the external AWS account
N. Update the route tables so that the AWS accounts can reach the ALB.
O. Choose a Network Load Balancer (NLB) as the type of load balancer for the ECS servic
P. Specify the NLB in the service definitio
Q. Create a VPC endpoint service for the NL
R. Share the VPC endpoint service with other AWS accounts.

Answer: D

NEW QUESTION 9
A company has created three VPCs: a production VPC, a nonproduction VPC, and a shared services VPC. The production VPC and the nonproduction VPC must each have communication with the shared services VPC. There must be no communication between the production VPC and the nonproduction VPC. A transit gateway is deployed to facilitate communication between VPCs.
Which route table configurations on the transit gateway will meet these requirements?

A. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for only the shared services VP
B. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.
C. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for each VP
D. Create an additional route table with only the shared services VPC attachment associated with propagated routes from each VPC.
E. Configure a route table with all the VPC attachments associated with propagated routes for only the shared services VPCCreate an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.
F. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes disable
G. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.

Answer: A

NEW QUESTION 10
A company is deploying a non-web application on an AWS load balancer. All targets are servers located
on-premises that can be accessed by using AWS Direct Connect. The company wants to ensure that the source IP addresses of clients connecting to the application are passed all the way to the end server.
How can this requirement be achieved?

A. Use a Network Load Balancer to automatically preserve the source IP address.
B. Use a Network Load Balancer and enable the X-Forwarded-For attribute.
C. Use a Network Load Balancer and enable the ProxyProtocol v2 attribute.
D. Use an Application Load Balancer to automatically preserve the source IP address in the X-Forwarded-For header.

Answer: C

Explanation:
https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#proxy-proto

NEW QUESTION 10
A company is using an AWS Site-to-Site VPN connection from the company's on-premises data center to a virtual private gateway in the AWS Cloud Because of congestion, the company is experiencing availability and performance issues as traffic travels across the internet before the traffic reaches AWS. A network engineer must reduce these issues for the connection as quickly as possible with minimum administration effort.
Which solution will meet these requirements?

A. Edit the existing Site-to-Site VPN connection by enabling acceleratio
B. Stop and start the VPN service on the customer gateway for the new setting to take effect.
C. Configure a transit gateway in the same AWS Region as the existing virtual private gatewa
D. Create a new accelerated Site-to-Site VPN connectio
E. Connect the new connection to the transit gateway by using a VPN attachmen
F. Update the customer gateway device to use the new Site to Site VPN connectio
G. Delete the existing Site-to-Site VPN connection
H. Create a new accelerated Site-to-Site VPN connectio
I. Connect the new Site-to-Site VPN connection to the existing virtual private gatewa
J. Update the customer gateway device to use the new Site-to-Site VPN connectio
K. Delete the existing Site-to-Site VPN connection.
L. Create a new AWS Direct Connect connection with a private VIF between the on-premises data center and the AWS Clou
M. Update the customer gateway device to use the new Direct Connect connectio
N. Delete the existing Site-to-Site VPN connection.

Answer: B

NEW QUESTION 11
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 ALB
D. and create a security rule to enforce forward secrecy (FS)
E. Change the ALB security policy to a policy that supports forward secrecy (FS)

Answer: D

NEW QUESTION 15
A global company operates all its non-production environments out of three AWS Regions: eu-west-1,
us-east-1, and us-west-1. The company hosts all its production workloads in two on-premises data centers. The company has 60 AWS accounts and each account has two VPCs in each Region. Each VPC has a virtual private gateway where two VPN connections terminate for resilient connectivity to the data centers. The company has 360 VPN tunnels to each data center, resulting in high management overhead. The total VPN throughput for each Region is 500 Mbps.
The company wants to migrate the production environments to AWS. The company needs a solution that will simplify the network architecture and allow for future growth. The production environments will generate an additional 2 Gbps of traffic per Region back to the data centers. This traffic will increase over time.
Which solution will meet these requirements?

A. Set up an AWS Direct Connect connection from each data center to AWS in each Regio
B. Create and attach private VIFs to a single Direct Connect gatewa
C. Attach the Direct Connect gateway to all the VPC
D. Remove the existing VPN connections that are attached directly to the virtual private gateways.
E. Create a single transit gateway with VPN connections from each data cente
F. Share the transit gateway with each account by using AWS Resource Access Manager (AWS RAM). Attach the transit gateway to each VP
G. Remove the existing VPN connections that are attached directly to the virtual private gateways.
H. Create a transit gateway in each Region with multiple newly commissioned VPN connections from each data cente
I. Share the transit gateways with each account by using AWS Resource Access Manager (AWS RAM). In each Region, attach the transit gateway to each
VPRemove the existing VPN connections that are attached directly to the virtual private gateways.
J. Peer all the VPCs in each Region to a new VPC in each Region that will function as a centralized transit VP
K. Create new VPN connections from each data center to the transit VPC
L. Terminate the original VPN connections that are attached to all the original VPC
M. Retain the new VPN connection to the new transit VPC in each Region.

Answer: C

NEW QUESTION 17
An organization is using a VPC endpoint for Amazon S3. When the security group rules for a set of instances were initially configured, access was restricted to allow traffic only to the IP addresses of the Amazon S3 API endpoints in the region from the published JSON file. The application was working properly, but now is logging a growing number of timeouts when connecting with Amazon S3. No internet gateway is configured for the VPC.
Which solution will fix the connectivity failures with the LEAST amount of effort?

A. Create a Lambda function to update the security group based on AmazonIPSpaceChanged notifications.
B. Update the VPC routing to direct Amazon S3 prefix-list traffic to the VPC endpoint using the route table APIs.
C. Update the application server's outbound security group to use the prefix-list for Amazon S3 in the same region.
D. Create an additional VPC endpoint for Amazon S3 in the same route table to scale the concurrent connections to Amazon.

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/subscribe-to-aws-public-ip-address-changes-via-amazon-sns/

NEW QUESTION 18
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 gateway in the VP
B. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT gateway.
C. Create an internet gateway and a NAT instance in the VP
D. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT instance.
E. Create an egress-only Internet gateway in the VPAdd a route to the existing subnet route tables topoint IPv6 traffic to the egress-only internet gateway.
F. Create an egress-only internet gateway in the VP
G. Configure a security group that denies all inbound traffi
H. Associate the security group with the egress-only internet gateway.

Answer: C

NEW QUESTION 19

A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin in an Amazon CloudFront distribution. The company wants to implement a custom authentication system that will provide a token for its authenticated customers.
The web application must ensure that the GET/POST requests come from authenticated customers before it delivers the content. A network engineer must design a solution that gives the web application the ability to identify authorized customers.
What is the MOST operationally efficient solution that meets these requirements?

A. Use the ALB to inspect the authorized token inside the GET/POST request payloa
B. Use an AWS Lambda function to insert a customized header to inform the web application of an authenticated customer request.
C. Integrate AWS WAF with the ALB to inspect the authorized token inside the GET/POST request payloa
D. Configure the ALB listener to insert a customized header to inform the web application of an authenticated customer request.
E. Use an AWS Lambda@Edge function to inspect the authorized token inside the GET/POST request payloa
F. Use the Lambda@Edge function also to insert a customized header to inform the web application of an authenticated customer request.
G. Set up an EC2 instance that has a third-party packet inspection tool to inspect the authorized token inside the GET/POST request payloa
H. Configure the tool to insert a customized header to inform the web application of an authenticated customer request.

Answer: C

NEW QUESTION 20
An international company provides early warning about tsunamis. The company plans to use IoT devices to monitor sea waves around the world. The data that is collected by the IoT devices must reach the company's infrastructure on AWS as quickly as possible. The company is using three operation centers around the world. Each operation center is connected to AWS through Its own AWS Direct Connect connection. Each operation center is connected to the internet through at least two upstream internet service providers.
The company has its own provider-independent (PI) address space. The IoT devices use TCP protocols for reliable transmission of the data they collect. The IoT devices have both landline and mobile internet connectivity. The infrastructure and the solution will be deployed in multiple AWS Regions. The company will use Amazon Route 53 for DNS services.
A network engineer needs to design connectivity between the IoT devices and the services that run in the AWS Cloud.
Which solution will meet these requirements with the HIGHEST availability?

A. Set up an Amazon CloudFront distribution with origin failove
B. Create an origin group for each Region where the solution is deployed.
C. Set up Route 53 latency-based routin
D. Add latency alias record
E. For the latency alias records, set the value of Evaluate Target Health to Yes.
F. Set up an accelerator in AWS Global Accelerato
G. Configure Regional endpoint groups andhealth checks.
H. Set up Bring Your Own IP (BYOIP) addresse
I. Use the same PI addresses for each Region where the solution is deployed.

Answer: B

Explanation:
https://aws.amazon.com/blogs/iot/automate-global-device-provisioning-with-aws-iot-core-and-amazon-route-53

NEW QUESTION 22
A network engineer is designing a hybrid architecture that uses a 1 Gbps AWS Direct Connect connection between the company's data center and two AWS Regions: us-east-1 and eu-west-1. The VPCs in us-east-1 are connected by a transit gateway and need to access several on-premises databases. According to company policy, only one VPC in eu-west-1 can be connected to one on-premises server. The on-premises network segments the traffic between the databases and the server.
How should the network engineer set up the Direct Connect connection to meet these requirements?

A. Create one hosted connectio
B. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use one Direc
C. Connect gateway for both VIFs to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.
D. Create one hosted connectio
E. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use two Direct Connect gateways, one for each VIF, to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.
F. Create one dedicated connectio
G. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use one Direct Connect gateway for both VIFs to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.
H. Create one dedicated connectio
I. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use two Direct Connect gateways, one for each VIF, to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.

Answer: B

Explanation:
This solution meets the requirements of the company by using a single Direct Connect connection with two VIFs, one connected to the transit gateway in us-east-1 and the other connected to the VPC in eu-west-1. Two Direct Connect gateways are used, one for each VIF, to route traffic from the Direct Connect location to the corresponding AWS Region along the path that has the lowest latency. This setup ensures that traffic between the VPCs in us-east-1 and on-premises databases is routed through the transit gateway, while traffic between the VPC in eu-west-1 and the on-premises server is routed directly through the private VIF.

NEW QUESTION 26
A company is developing an application in which IoT devices will report measurements to the AWS Cloud. The application will have millions of end users. The company observes that the IoT devices cannot support DNS resolution. The company needs to implement an Amazon EC2 Auto Scaling 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 grou
B. Attach the Auto Scaling group to the AL
C. Set up the IoT devices to connect to the IP addresses of the NLB.

D. Use an AWS Global Accelerator accelerator with an Application Load Balancer (ALB) endpoin
E. Create an EC2 Auto Scaling grou
F. Attach the Auto Scaling group to the ALSet up the IoT devices to connect to the IP addresses of the accelerator.
G. Use a Network Load Balancer (NLB). Create an EC2 Auto Scaling grou
H. Attach the Auto Scaling group to the NL
I. Set up the IoT devices to connect to the IP addresses of the NLB.
J. Use an AWS Global Accelerator accelerator with a Network Load Balancer (NLB) endpoin
K. Create anEC2 Auto Scaling grou
L. Attach the Auto Scaling group to the NL
M. Set up the IoT devices to connect to the IP addresses of the accelerator.

Answer: D

Explanation:
AWS Global Accelerator can provide static IP addresses that the IoT devices can connect to without using DNS2. It can also route traffic over the AWS global network and improve performance and availability for the IoT devices2. An NLB can provide end-to-end encryption for HTTPS traffic by using TLS as a target group protocol and terminating SSL connections at the load balancer level1. An NLB can also support session affinity (sticky sessions) with TCP connections1.

NEW QUESTION 27
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 VPC
B. Use the new connection to connect additional VPCs.
C. Create virtual private gateways for each VPC that is over the service quot
D. Use AWS Site-to-Site VPNto connect the virtual private gateways to the corporate network.
E. Create a Direct Connect gateway, and add virtual private gateway associations to the VPC
F. Configure a private VIF to connect to the corporate network.
G. Create a transit gateway, and attach the VPC
H. Create a Direct Connect gateway, and associate it with the transit gatewa
I. Create a transit VIF to the Direct Connect gateway.

Answer: D

Explanation:
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 transitgateway, 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 29
An ecommerce company is hosting a web application on Amazon EC2 instances to handle continuously changing customer demand. The EC2 instances are part of an Auto Scaling group. The company wants to implement a solution to distribute traffic from customers to the EC2 instances. The company must encrypt all traffic at all stages between the customers and the application servers. No decryption at intermediate points is allowed.
Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB). Add an HTTPS listener to the AL
B. Configure the Auto Scaling group to register instances with the ALB's target group.
C. Create an Amazon CloudFront distributio
D. Configure the distribution with a custom SSL/TLS certificat
E. Set the Auto Scaling group as the distribution's origin.
F. Create a Network Load Balancer (NLB). Add a TCP listener to the NL
G. Configure the Auto Scaling group to register instances with the NLB's target group.
H. Create a Gateway Load Balancer (GLB). Configure the Auto Scaling group to register instances with the GLB's target group.

Answer: C

Explanation:
To distribute traffic from customers to EC2 instances in an Auto Scaling group and encrypt all traffic at all stages between the customers and the application servers without decryption at intermediate points, the company should create a Network Load Balancer (NLB) with a TCP listener and configure the Auto Scaling group to register instances with the NLB's target group (Option C). This solution allows for end-to-end encryption of traffic without decryption at intermediate points.

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: DEF

Explanation:
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 35
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

ANS-C01 Practice Exam Features:

* ANS-C01 Questions and Answers Updated Frequently

* ANS-C01 Practice Questions Verified by Expert Senior Certified Staff

* ANS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* ANS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The ANS-C01 Practice Test Here