AWS-Certified-Advanced-Networking-Specialty Dumps

Amazon AWS Certified Advanced Networking - Specialty

https://www.certleader.com/AWS-Certified-Advanced-Networking-Specialty-dumps.html

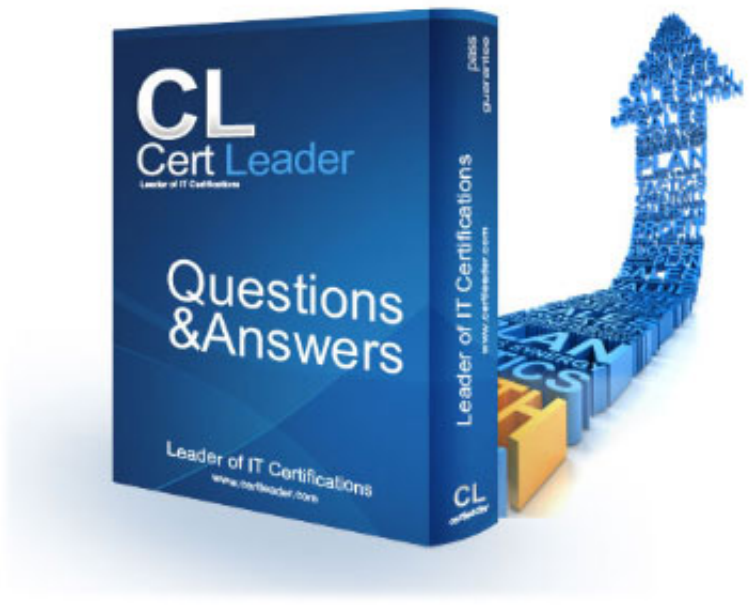

NEW QUESTION 1
A retail company is running its service on AWS. The company's architecture includes Application Load Balancers (ALBs) in public subnets. The ALB target groups are configured to send traffic to backend Amazon EC2 instances in private subnets. These backend EC2 instances can call externally hosted services over the internet by using a NAT gateway.
The company has noticed in its billing that NAT gateway usage has increased significantly. A network engineer needs to find out the source of this increased usage.
Which options can the network engineer use to investigate the traffic through the NAT gateway? (Choose two.)

A. Enable VPC flow logs on the NAT gateway's elastic network interfac
B. Publish the logs to a log group in Amazon CloudWatch Log
C. Use CloudWatch Logs Insights to query and analyze the logs.
D. Enable NAT gateway access log
E. Publish the logs to a log group in Amazon CloudWatch Log
F. Use CloudWatch Logs Insights to query and analyze the logs.
G. Configure Traffic Mirroring on the NAT gateway's elastic network interfac
H. Send the traffic to an additional EC2 instanc
I. Use tools such as tcpdump and Wireshark to query and analyze the mirrored traffic.
J. Enable VPC flow logs on the NAT gateway's elastic network interfac
K. Publish the logs to an Amazon S3 bucke
L. Create a custom table for the S3 bucket in Amazon Athena to describe the log structur
M. Use Athena to query and analyze the logs.
N. Enable NAT gateway access log
O. Publish the logs to an Amazon S3 bucke
P. Create a custom table for the S3 bucket in Amazon Athena to describe the log structur
Q. Use Athena to query and analyze the logs.

Answer: AD

Explanation:
To investigate the increased usage of a NAT gateway in a VPC architecture with ALBs and backend EC2 instances, a network engineer can use the following options:
- Enable VPC flow logs on the NAT gateway's elastic network interface and publish the logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query and analyze the logs.
(Option A)
- Enable VPC flow logs on the NAT gateway's elastic network interface and publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure and use Athena to query and analyze the logs. (Option D)

These options allow for detailed analysis of traffic through the NAT gateway to identify the source of increased usage.

NEW QUESTION 2
A government contractor is designing a multi-account environment with multiple VPCs for a customer. A network security policy requires all traffic between any two VPCs to be transparently inspected by a
third-party appliance.
The customer wants a solution that features AWS Transit Gateway. The setup must be highly available across multiple Availability Zones, and the solution needs to support automated failover. Furthermore, asymmetric routing is not supported by the inspection appliances.
Which combination of steps is part of a solution that meets these requirements? (Choose two.)

A. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VP
B. Connect the inspection VPC to the transit gateway by using a VPCattachmen
C. Create a target group, and register the appliances with the target grou
D. Create a Network Load Balancer (NLB), and set it up to forward to the newly created target grou
E. Configure a default route in the inspection VPCs transit gateway subnet toward the NLB.
F. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VP
G. Connect the inspection VPC to the transit gateway by using a VPC attachmen
H. Create a target group, and register the appliances with the target grou
I. Create a Gateway Load Balancer, and set it up to forward to the newly created target grou
J. Configure a default route in the inspection VPC's transit gateway subnet toward the Gateway Load Balancer endpoint.
K. Configure two route tables on the transit gatewa
L. Associate one route table with all the attachments of the application VPC
M. Associate the other route table with the inspection VPC's attachmen
N. Propagate all VPC attachments into the inspection route tabl
O. Define a static default route in the application route tabl
P. Enable appliance mode on the attachment that connects the inspection VPC.
Q. Configure two route tables on the transit gatewa
R. Associate one route table with all the attachments of the application VPC
S. Associate the other route table with the inspection VPCs attachmen
T. Propagate all VPC attachments into the application route tabl
. Define a static default route in the inspection route tabl
. Enable appliance mode on the attachment that connects the inspection VPC.
. Configure one route table on the transit gatewa
. Associate the route table with all the VPC
. Propagate all VPC attachments into the route tabl
. Define a static default route in the route table.

Answer: BC

NEW QUESTION 3
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 4
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 5
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 6
An IoT company sells hardware sensor modules that periodically send out temperature, humidity, pressure, and location data through the MQTT messaging protocol. The hardware sensor modules send this data to the company's on-premises MQTT brokers that run on Linux servers behind a load balancer. The hardware sensor modules have been hardcoded with public IP addresses to reach the brokers.
The company is growing and is acquiring customers across the world. The existing solution can no longer scale and is introducing additional latency because of the company's global presence. As a result, the company decides to migrate its entire infrastructure from on premises to the AWS Cloud. The company needs to migrate without reconfiguring the hardware sensor modules that are already deployed across the world. The solution also must minimize latency.
The company migrates the MQTT brokers to run on Amazon EC2 instances. What should the company do next to meet these requirements?

A. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listener
B. Use Bring Your Own IP (BYOIP) from the on-premises network with the NLB.
C. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listener
D. Create an AWS Global Accelerator accelerator in front of the NLUse Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator.
E. Place the EC2 instances behind an Application Load Balancer (ALB). Configure TCP listener
F. Create an AWS Global Accelerator accelerator in front of the AL
G. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator
H. Place the EC2 instances behind an Amazon CloudFront distributio
I. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

Answer: B

NEW QUESTION 7
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 thatfilters 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. Create a VPN connection over the Direct Connect connection by using the on-premises firewal
B. Use the firewall to block all traffic from on premises to AW
C. Allow a stateful connection from the EC2 instances to initiate the requests.
D. Configure the on-premises firewall to filter all requests from the on-premises network to the EC2 instance
E. Allow a stateful connection if the EC2 instances in the VPC initiate the traffic.
F. Deploy a NAT gateway into a private subnet in the VPC where the EC2 instances are deploye
G. Specify the NAT gateway type as privat
H. Configure the on-premises firewall to allow connections from the IP address that is assigned to the NAT gateway.
I. 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.

Answer: C

NEW QUESTION 8
A company has an AWS Site-to-Site VPN connection between its existing VPC and on-premises network. The default DHCP options set is associated with the VPC. The company has an application that is running on an Amazon Linux 2 Amazon EC2 instance in the VPC. The application must retrieve an Amazon RDS database secret that is stored in AWS Secrets Manager through a private VPC endpoint. An on-premises application provides internal RESTful API service that can be reached by URL (https://api.example.internal). Two on-premises Windows DNS servers provide internal DNS resolution.
The application on the EC2 instance needs to call the internal API service that is deployed in the on-premises environment. When the application on the EC2 instance attempts to call the internal API service by referring to the hostname that is assigned to the service, the call fails. When a network engineer tests the API service call from the same EC2 instance by using the API service's IP address, the call is successful.
What should the network engineer do to resolve this issue and prevent the same problem from affecting other resources in the VPC?

A. Create a new DHCP options set that specifies the on-premises Windows DNS server
B. Associate the new DHCP options set with the existing VP
C. Reboot the Amazon Linux 2 EC2 instance.
D. Create an Amazon Route 53 Resolver rul
E. Associate the rule with the VP
F. Configure the rule to forward DNS queries to the on-premises Windows DNS servers if the domain name matches example.internal.
G. Modify the local host file in the Amazon Linux 2 EC2 instance in the VPMap the service domain name (api.example.internal) to the IP address of the internal API service.
H. Modify the local /etc/resolv.conf file in the Amazon Linux 2 EC2 instance in the VP
I. Change the IP addresses of the name servers in the file to the IP addresses of the company's on-premises Windows DNS servers.

Answer: B

Explanation:
Creating an Amazon Route 53 Resolver rule and associating it with the VPC would enable forwarding of DNS queries for a specified domain name (example.internal) to a specified IP address (the on-premises Windows DNS servers)3. This would allow EC2 instances in the VPC to resolve the internal API service by using its hostname. Configuring the rule to forward DNS queries only if the domain name matches example.internal would also allow EC2 instances to use the Amazon Route 53 Resolver server for other DNS queries, such as those for AWS services through private VPC endpoints2.

NEW QUESTION 9
A data analytics company has a 100-node high performance computing (HPC) cluster. The HPC cluster is for parallel data processing and is hosted in a VPC in the AWS Cloud. As part of the data processing workflow, the HPC cluster needs to perform several DNS queries to resolve and connect to Amazon RDS databases, Amazon S3 buckets, and on-premises data stores that are accessible through AWS Direct Connect. The HPC cluster can increase in size by five to seven times during the company’s peak event at the end of the year.
The company is using two Amazon EC2 instances as primary DNS servers for the VPC. The EC2 instances are configured to forward queries to the default VPC resolver for Amazon Route 53 hosted domains and to the on-premises DNS servers for other on-premises hosted domain names. The company notices job

failures and finds that DNS queries from the HPC cluster nodes failed when the nodes tried to resolve RDS and S3 bucket endpoints.
Which architectural change should a network engineer implement to provide the DNS service in the MOST scalable way?

A. Scale out the DNS service by adding two additional EC2 instances in the VP
B. Reconfigure half of the HPC cluster nodes to use these new DNS server
C. Plan to scale out by adding additional EC2instance-based DNS servers in the future as the HPC cluster size grows.
D. Scale up the existing EC2 instances that the company is using as DNS server
E. Change the instance size to the largest possible instance size to accommodate the current DNS load and theanticipated load in the future.
F. Create Route 53 Resolver outbound endpoint
G. Create Route 53 Resolver rules to forward queries to on-premises DNS servers for on premises hosted domain name
H. Reconfigure the HPC cluster nodes to use the default VPC resolver instead of the EC2 instance-based DNS server
I. Terminate the EC2 instances.
J. Create Route 53 Resolver inbound endpoint
K. Create rules on the on-premises DNS servers to forward queries to the default VPC resolve
L. Reconfigure the HPC cluster nodes to forward all DNS queries to the on-premises DNS server
M. Terminate the EC2 instances.

Answer: C

NEW QUESTION 10
A company has an AWS Direct Connect connection between its on-premises data center in the United States (US) and workloads in the us-east-1 Region. The connection uses a transit VIF to connect the data center to a transit gateway in us-east-1.
The company is opening a new office in Europe with a new on-premises data center in England. A Direct Connect connection will connect the new data center with some workloads that are running in a single VPC in the eu-west-2 Region. The company needs to connect the US data center and us-east-1 with the Europe data center and eu-west-2. A network engineer must establish full connectivity between the data centers and Regions with the lowest possible latency.
How should the network engineer design the network architecture to meet these requirements?

A. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VI
B. Associate the transit gateway in us-east-1 with the same Direct Connect gatewa
C. Enable SiteLink for the transit VIF and the private VIF.
D. Connect the VPC in eu-west-2 to a new transit gatewa
E. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VI
F. Associate the transit gateway in us-east-1 with the same Direct Connect gatewa
G. Enable SiteLink for both transit VIF
H. Peer the two transit gateways.
I. Connect the VPC in eu-west-2 to a new transit gatewa
J. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VI
K. Create a new Direct Connect gatewa
L. Associate the transit gateway in us-east-1 with the new Direct Connect gatewa
M. Enable SiteLink for both transit VIF
N. Peer the two transit gateways.
O. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VI
P. Create a new Direct Connect gatewa
Q. Associate the transit gateway in us-east-1 with the new Direct Connect gatewa
R. Enable SiteLink for the transit VIF and the private VIF.

Answer: C

NEW QUESTION 10
A company has its production VPC (VPC-A) in the eu-west-1 Region in Account 1. VPC-A is attached to a transit gateway (TGW-A) that is connected to an on-premises data center in Dublin, Ireland, by an AWS
Direct Connect transit VIF that is configured for an AWS Direct Connect gateway. The company also has a staging VPC (VPC-B) that is attached to another transit gateway (TGW-B) in the eu-west-2 Region in Account 2.
A network engineer must implement connectivity between VPC-B and the on-premises data center in Dublin. Which solutions will meet these requirements?
(Choose two.)

A. Configure inter-Region VPC peering between VPC-A and VPC-
B. Add the required VPC peering route
C. Add the VPC-B CIDR block in the allowed prefixes on the Direct Connect gateway association.
D. Associate TGW-B with the Direct Connect gatewa
E. Advertise the VPC-B CIDR block under the allowed prefixes.
F. Configure another transit VIF on the Direct Connect connection and associate TGW-
G. Advertise the VPC-B CIDR block under the allowed prefixes.
H. Configure inter-Region transit gateway peering between TGW-A and TGW-
I. Add the peering routes in the transit gateway route table
J. Add both the VPC-A and the VPC-B CIDR block under the allowed prefix list in the Direct Connect gateway association.
K. Configure an AWS Site-to-Site VPN connection over the transit VIF to TGW-B as a VPN attachment.

Answer: BC

Explanation:
* B. Associate TGW-B with the Direct Connect gateway. Advertise the VPC-B CIDR block under the allowed prefixes. This will allow traffic from VPC-B to be sent over the Direct Connect connection to the on-premises data center via TGW-B. C. Configure another transit VIF on the Direct Connect connection and associate TGW-B. Advertise the VPC-B CIDR block under the allowed prefixes. This will enable the use of the Direct Connect connection for VPC-B's traffic by connecting TGW-B to the Direct Connect gateway.

NEW QUESTION 12
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 14
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 18
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 gatewa
B. Provide the Connectivity account I
C. 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 attachmen
D. Associate a route table with the attachment.
E. * 1. In the Production account: Create a resource share in AWS Resource Access Manager for the VPC subnet
F. Provide the Connectivity account I
G. 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 attachmen
H. Associate a route table with the attachment.
I. * 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the VPC subnet
J. Provide the Production account I
K. 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 attachmen
L. Associate a route table with the attachment.
M. * 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the transit gatewa
N. 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 attachmen
O. Associate a route table with the attachment.

Answer: A

Explanation:
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 22
A customer has set up multiple VPCs for Dev, Test, Prod, and Management. You need to set up AWS Direct Connect to enable data flow from on-premises to each VPC. The customer hasmonitoring software running in the Management VPC that collects metrics from the instances in all the other VPCs. Due to budget requirements, data transfer charges should be kept at minimum.
Which design should be recommended?

A. Create a total of four private VIFs, one for each VPC owned by the customer, and route traffic between VPCs using the Direct Connect link.
B. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs.
C. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs, enable source/destination NAT in the Management VPC.
D. Create a total of four private VIFs, and enable VPC peering between all VPCs.

Answer: D

Explanation:
- creating VPC peering is free of charge - traffic costs ~0.01€/GB for VPC peering (IN + OUT) and ~0.02€/GB for direct connect (OUT only). As the communication involved in monitoring will never have IN == OUT, then 0.01 * (IN + OUT) will always be lower the 0.02 * OUT, ergo VPC peering will be cheaper

NEW QUESTION 26
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. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering addres
B. Create a new VPN connection that supports IPv6 connectivit
C. Add an egress-only internet gatewa
D. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices
E. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering addres
F. Update the existing VPN connection to support IPv6 connectivit
G. Add an egress-only internet gatewa
H. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
I. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering addres
J. Create a new VPN connection that supports IPv6 connectivit
K. Add an egress-only internet gatewa
L. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
M. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering addres
N. Create a new VPN connection that supports IPv6 connectivit
O. Add a NAT gatewa
P. 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: B

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

Explanation:
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 30
A global company runs business applications in the us-east-1 Region inside a VPC. One of the company's regional offices in London uses a virtual private gateway for an AWS Site-to-Site VPN connection tom the VPC. The company has configured a transit gateway and has set up peering between the VPC and other VPCs that various departments in the company use.
Employees at the London office are experiencing latency issues when they connect to the business applications.
What should a network engineer do to reduce this latency?

A. Create a new Site-to-Site VPN connectio
B. Set the transit gateway as the target gatewa
C. Enable acceleration on the new Site-to-Site VPN connectio
D. Update the VPN device in the London office with the new connection details.

E. Modify the existing Site-to-Site VPN connection by setting the transit gateway as the target gateway.Enable acceleration on the existing Site-to-Site VPN connection.
F. Create a new transit gateway in the eu-west-2 (London) Regio
G. Peer the new transit gateway with the existing transit gatewa
H. Modify the existing Site-to-Site VPN connection by setting the new transit gateway as the target gateway.
I. Create a new AWS Global Accelerator standard accelerator that has an endpoint of the Site-to-Site VPN connectio
J. Update the VPN device in the London office with the new connection details.

Answer: A

Explanation:
Enabling acceleration for a Site-to-Site VPN connection uses AWS Global Accelerator to route traffic from the on-premises network to an AWS edge location that is closest to the customer gateway device1. AWS Global Accelerator optimizes the network path, using the congestion-free AWS global network to route traffic to the endpoint that provides the best application performance2. Setting the transit gateway as the target gateway enables connectivity between the on-premises network and multiple VPCs that are attached to the transit gateway3.

NEW QUESTION 32
An AWS CloudFormation template is being used to create a VPC peering connection between two existing operational VPCs, each belonging to a different AWS account. All necessary components in the 'Remote' (receiving) account are already in place.
The template below creates the VPC peering connection in the Originating account. It contains these components:
AWSTemplateFormation Version: 2010-09-09 Parameters:
Originating VCId: Type: String RemoteVPCId: Type: String
RemoteVPCAccountId: Type: String Resources:
newVPCPeeringConnection:
Type: 'AWS::EC2::VPCPeeringConnection'
Properties:
VpcdId: !Ref OriginatingVPCId PeerVpcId: !Ref RemoteVPCId PeerOwnerId: !Ref RemoteVPCAccountId
Which additional AWS CloudFormation components are necessary in the Originating account to create an operational cross-account VPC peering connection with AWS CloudFormation? (Select two.)

A. Resources:NewEC2SecurityGroup:Type: AWS::EC2::SecurityGroup
B. Resources:NetworkInterfaceToRemoteVPC:Type: "AWS::EC2NetworkInterface"
C. Resources:newEC2Route:Type: AWS::EC2::Route
D. Resources:VPCGatewayToRemoteVPC:Type: "AWS::EC2::VPCGatewayAttachment"
E. Resources:newVPCPeeringConnection:Type: 'AWS::EC2VPCPeeringConnection'PeerRoleArn: !Ref PeerRoleArn

Answer: CE

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/AWS_EC2.html

NEW QUESTION 34
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 subnet
B. Create an S3 interface endpoint in the VP
C. Modify the application configuration to use the S3 endpoint-specific DNS hostname.
D. Deploy the EC2 instances in the private subnet
E. Create a NAT gateway in the VP
F. Create default routes in the private subnets to the NAT gatewa
G. Connect to Amazon S3 by using the NAT gateway.
H. Deploy the EC2 instances in the private subnet
I. Create an S3 gateway endpoint in the VPSpecify die route table of the private subnets during endpoint creation to create routes to Amazon S3.
J. Deploy the EC2 instances in the private subnet
K. Create an S3 interface endpoint in the VP
L. Modify the application configuration to use the S3 endpoint-specific DNS hostname.

Answer: C

Explanation:
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 35
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 VPCreate 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 36
A company hosts an application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company recently experienced a network security breach. A network engineer must collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application.
What is the MOST operationally efficient solution that meets these requirements?

A. Configure the ALB to store logs in an Amazon S3 bucke
B. Download the files from Amazon S3, and use a spreadsheet application to analyze the logs.
C. Configure the ALB to push logs to Amazon Kinesis Data Stream
D. Use Amazon Kinesis Data Analytics to analyze the logs.
E. Configure Amazon Kinesis Data Streams to stream data from the ALB to Amazon OpenSearch Service (Amazon Elasticsearch Service). Use search operations in Amazon OpenSearch Service (Amazon Elasticsearch Service) to analyze the data.
F. Configure the ALB to store logs in an Amazon S3 bucke
G. Use Amazon Athena to analyze the logs in Amazon S3.

Answer: D

Explanation:
The most operationally efficient solution to collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application would be to configure the ALB to store logs in an Amazon S3 bucket and use Amazon Athena to analyze the logs in Amazon S3 (Option D). This solution allows for quick and easy analysis of log data without requiring manual download or manipulation of log files.

NEW QUESTION 37
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 42
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 Gatewa
E. Attach the transit gateway to the VPC, and connect the Direct Connect gateway to the transit gateway.

Answer: B

Explanation:
"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 45
A company has deployed its AWS environment in a single AWS Region. The environment consists of a few hundred application VPCs, a shared services VPC, and a VPN connection to the company's on-premises environment. A network engineer needs to implement a transit gateway with the following requirements:
• Application VPCs must be isolated from each other.
• Bidirectional communication must be allowed between the application VPCs and the on-premises network.
• Bidirectional communication must be allowed between the application VPCs and the shared services VPC. The network engineer creates the transit gateway with options disabled for default route table association and default route table propagation. The network engineer also creates the VPN attachment for the on-premises network and creates the VPC attachments for the application VPCs and the shared services VPC.
The network engineer must meet all the requirements for the transit gateway by designing a solution that needs the least number of transit gateway route tables.
Which combination of actions should the network engineer perform to accomplish this goal?(Choose two.)

A. Configure a separate transit gateway route table for on premise
B. Associate the VPN attachment with this transit gateway route tabl
C. Propagate all application VPC attachments to this transit gateway route table.

D. Configure a separate transit gateway route table for each application VP
E. Associate each application VPC attachment with its respective transit gateway route tabl
F. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.
G. Configure a separate transit gateway route table for all application VPC
H. Associate all application VPCs with this transit gateway route tabl
I. Propagate the shared services VPC attachment and the VPNattachment to this transit gateway route table.
J. Configure a separate transit gateway route table for the shared services VP
K. Associate the shared services VPC attachment with this transit gateway route tabl
L. Propagate all application VPC attachments to this transit gateway route table.
M. Configure a separate transit gateway route table for on premises and the shared services VP
N. Associate the VPN attachment and the shared services VPC attachment with this transit gateway route tabl
O. Propagate all application VPC attachments to this transit gateway route table.

Answer: BD

NEW QUESTION 46
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 48
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 53
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 7224:7100. Set the BGP community tag for all prefixes from the failover data center to 7224:7300
B. 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
C. 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
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: B

NEW QUESTION 54
A company's network engineer builds and tests network designs for VPCs in a development account. The company needs to monitor the changes that are made to network resources and must ensure strict compliance with network security policies. The company also needs access to the historical configurations of network resources.
Which solution will meet these requirements?

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a custom pattern to monitor the account for change
B. Configure the rule to invoke an AWS Lambda function to identify noncompliant resource
C. Update an Amazon DynamoDB table with the changes that are identified.

D. Create custom metrics from Amazon CloudWatch log
E. Use the metrics to invoke an AWS Lambda function to identify noncompliant resource
F. Update an Amazon DynamoDB table with the changes that are identified.
G. Record the current state of network resources by using AWS Confi
H. Create rules that reflect the desired configuration setting
I. Set remediation for noncompliant resources.
J. Record the current state of network resources by using AWS Systems Manager Inventor
K. Use Systems Manager State Manager to enforce the desired configuration settings and to carry out remediation for noncompliant resources.

Answer: C

Explanation:
Recording the current state of network resources by using AWS Config would enable auditing and assessment of resource configurations and compliance3.
Creating rules that reflect the desired configuration settings would enable evaluation of whether the network resources comply with network security policies3.
Setting remediation for noncompliant resources would enable automatic correction of undesired configurations3.

NEW QUESTION 56
A bank built a new version of its banking application in AWS using containers that content to an on-premises database over VPN connection. This application version requires users to also update their client application. The bank plans to deprecate the earlier client version. However, the company wants to keep supporting earlier clients through their on-premises version of the application to serve a small portion of the customers who haven't yet upgraded.
What design will allow the company to serve both newer and earlier clients in the MOST efficient way?

A. Use an Amazon Route 53 multivalue answer routing policy to route older client traffic to the on-premises application version and the rest of the traffic to the new AWS based version.
B. Use a Classic Load Balancer for the new applicatio
C. Route all traffic to the new application by using an Elastic Load Balancing (ELB) load balancer DN
D. Define a user-agent-based rule on the backend servers to redirect earlier clients to the on-premises application.
E. Use an Application Load Balancer for the new applicatio
F. Register both the new and earlier applications as separate target groups and use path-based routing to route traffic based on the application version.
G. Use an Application Load Balancer for the new applicatio
H. Register both the new and earlier application backends as separate target group
I. Use header-based routing to route traffic based on the application version.

Answer: D

NEW QUESTION 59
A company is deploying a new application on AWS. The application uses dynamic multicasting. The company has five VPCs that are all attached to a transit gateway Amazon EC2 instances in each VPC need to be able to register dynamically to receive a multicast transmission.
How should a network engineer configure the AWS resources to meet these requirements?

A. Create a static source multicast domain within the transit gatewa
B. Associate the VPCs and applicable subnets with the multicast domai
C. Register the multicast senders' network interface with the multicast domai
D. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
E. Create a static source multicast domain within the transit gatewa
F. Associate the VPCs and applicable subnets with the multicast domai
G. Register the multicast senders' network interface with the multicast domai
H. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.
I. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway.Associate the VPCs and applicable subnets with the multicast domai
J. Register the multicast senders' network interface with the multicast domai
K. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
L. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway.Associate the VPCs and applicable subnets with the multicast domai
M. Register the multicast senders' network interface with the multicast domai
N. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.

Answer: C

NEW QUESTION 60
A company's AWS architecture consists of several VPCs. The VPCs include a shared services VPC and several application VPCs. The company has established network connectivity from all VPCs to the
on-premises DNS servers.
Applications that are deployed in the application VPCs must be able to resolve DNS for internally hosted domains on premises. The applications also must be able to resolve local VPC domain names and domains that are hosted in Amazon Route 53 private hosted zones.
What should a network engineer do to meet these requirements?

A. Create a new Route 53 Resolver inbound endpoint in the shared services VP
B. Create forwarding rules for the on-premises hosted domain
C. Associate the rules with the new Resolver endpoint and each application VP
D. Update each application VPC's DHCP configuration to point DNS resolution to the new Resolver endpoint.
E. Create a new Route 53 Resolver outbound endpoint in the shared services VP
F. Create forwarding rules for the on-premises hosted domain
G. Associate the rules with the new Resolver endpoint and each application VPC.
H. Create a new Route 53 Resolver outbound endpoint in the shared services VPCreate forwarding rules for the on-premises hosted domain
I. Associate the rules with the new Resolver endpoint and each application VPUpdate each application VPC's DHCP configuration to point DNS resolution to the new Resolver endpoint.
J. Create a new Route 53 Resolver inbound endpoint in the shared services VP
K. Create forwarding rules for the on-premises hosted domain
L. Associate the rules with the new Resolver endpoint and each application VPC.

Answer: B

Explanation:
Creating a new Route 53 Resolver outbound endpoint in the shared services VPC would enable forwarding of DNS queries from the VPC to on-premises1. Creating forwarding rules for the on-premises hosted domains would enable specifying which domain names are forwarded to the on-premises DNS servers2. Associating the rules with the new Resolver endpoint and each application VPC would enable applying the rules to the VPCs2. This solution would not affect the default DNS resolution behavior of Route 53 Resolver for local VPC domain names and domains that are hosted in Route 53 private hosted zones3.

NEW QUESTION 61
A company has several production applications across different accounts in the AWS Cloud. The company operates from the us-east-1 Region only. Only certain partner companies can access the applications. The applications are running on Amazon EC2 instances that are in an Auto Scaling group behind an Application Load Balancer (ALB). The EC2 instances are in private subnets and allow traffic only from the ALB. The ALB is in a public subnet and allows inbound traffic only from partner network IP address ranges over port 80.
When the company adds a new partner, the company must allow the IP address range of the partner network in the security group that is associated with the ALB in each account. A network engineer must implement a solution to centrally manage the partner network IP address ranges.
Which solution will meet these requirements in the MOST operationally efficient manner?

A. Create an Amazon DynamoDB table to maintain all IP address ranges and security groups that need to be update
B. Update the DynamoDB table with the new IP address range when the company adds a new partne
C. Invoke an AWS Lambda function to read new IP address ranges and security groups from the DynamoDB table to update the security group
D. Deploy this solution in all accounts.
E. Create a new prefix lis
F. Add all allowed IP address ranges to the prefix lis
G. Use Amazon EventBridge (Amazon CloudWatch Events) rules to invoke an AWS Lambda function to update security groups whenever a new IP address range is added to the prefix lis
H. Deploy this solution in all accounts.
I. Create a new prefix lis
J. Add all allowed IP address ranges to the prefix lis
K. Share the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM). Update security groups to use the prefix list instead of the partner IP address rang
L. Update the prefix list with the new IP address range when the company adds a new partner.
M. Create an Amazon S3 bucket to maintain all IP address ranges and security groups that need to be update
N. Update the S3 bucket with the new IP address range when the company adds a new partne
O. Invoke an AWS Lambda function to read new IP address ranges and security groups from the S3 bucket to update the security group
P. Deploy this solution in all accounts.

Answer: C

Explanation:
Creating a new prefix list and adding all allowed IP address ranges to the prefix list would enable grouping of CIDR blocks that can be referenced in security group rules3. Sharing the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM)would enable central management of the partner network IP address ranges5. Updating security groups to use the prefix list instead of the partner IP address range would enable simplification of security group rules3. Updating the prefix list with the new IP address range when the company adds a new partner would enable automatic propagation of the changes to all security groups that use the prefix list3.

NEW QUESTION 66
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Advanced-Networking-Specialty Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Certified-Advanced-Networking-Specialty-dumps.html