Amazon-Web-Services

Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam

NEW QUESTION 1
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 2
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 3
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 4
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 5
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 6
A real estate company is building an internal application so that real estate agents can upload photos and videos of various properties. The application will store these photos and videos in an Amazon S3 bucket as objects and will use Amazon DynamoDB to store corresponding metadata. The S3 bucket will be configured to publish all PUT events for new object uploads to an Amazon Simple Queue Service (Amazon SQS) queue.
A compute cluster of Amazon EC2 instances will poll the SQS queue to find out about newly uploaded objects. The cluster will retrieve new objects, perform proprietary image and video recognition and classification update metadata in DynamoDB and replace the objects with new watermarked objects. The company does not want public IP addresses on the EC2 instances.
Which networking design solution will meet these requirements MOST cost-effectively as application usage increases?

A. Place the EC2 instances in a public subne
B. Disable the Auto-assign Public IP option while launching the EC2 instance
C. Create an internet gatewa
D. Attach the internet gateway to the VP
E. In the public subnet's route table, add a default route that points to the internet gateway.
F. Place the EC2 instances in a private subne
G. Create a NAT gateway in a public subnet in the same Availability Zon
H. Create an internet gatewa
I. Attach the internet gateway to the VP
J. In the public subnet's route table, add a default route that points to the internet gateway
K. Place the EC2 instances in a private subne
L. Create an interface VPC endpoint for Amazon SQ
M. Create gateway VPC endpoints for Amazon S3 and DynamoDB.
N. Place the EC2 instances in a private subne
O. Create a gateway VPC endpoint for Amazon SQS.Create interface VPC endpoints for Amazon S3 and DynamoDB.

Answer: C

NEW QUESTION 7
All IP addresses within a 10.0.0.0/16 VPC are fully utilized with application servers across two Availability Zones. The application servers need to send frequent UDP probes to a single central authentication server on the Internet to confirm that is running up-to-date packages. The network is designed for application servers to use a single NAT gateway for internal access. Testing reveals that a few of the servers are unable to communicate with the authentication server.

A. The NAT gateway does not support UDP traffic.
B. The authentication server is not accepting traffic.
C. The NAT gateway cannot allocate more ports.

D. The NAT gateway is launched in a private subnet.

Answer: C

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

NEW QUESTION 8
A company is deploying third-party firewall appliances for traffic inspection and NAT capabilities in its VPC. The VPC is configured with private subnets and public subnets. The company needs to deploy the firewall appliances behind a load balancer.
Which architecture will meet these requirements MOST cost-effectively?

A. Deploy a Gateway Load Balancer with the firewall appliances as target
B. Configure the firewall appliances with a single network interface in a private subne
C. Use a NAT gateway to send the traffic to the internet after inspection.
D. Deploy a Gateway Load Balancer with the firewall appliances as target
E. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subne
F. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.
G. Deploy a Network Load Balancer with the firewall appliances as target
H. Configure the firewall appliances with a single network interface in a private subne
I. Use a NAT gateway to send the traffic to the internet after inspection.
J. Deploy a Network Load Balancer with the firewall appliances as target
K. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subne
L. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.

Answer: B

NEW QUESTION 9
Your organization has a newly installed 1-Gbps AWS Direct Connect connection. You order the cross-connect from the Direct Connect location provider to the port on your router in the same facility. To enable the use of your first virtual interface, your router must be configured appropriately.
What are the minimum requirements for your router?

A. 1-Gbps Multi Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
B. 1-Gbps Single Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
C. IPsec Parameters, Pre-Shared key, Peer IP Address, BGP Session with MD5
D. BGP Session with MD5, 802.1Q VLAN, Route-Map, Prefix List, IPsec encrypted GRE Tunnel

Answer: B

NEW QUESTION 10
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 10
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 11
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 14
A company is planning a migration of its critical workloads from an on-premises data center to Amazon EC2 instances. The plan includes a new 10 Gbps AWS Direct Connect dedicated connection from the on-premises data center to a VPC that is attached to a transit gateway. The migration must occur over encrypted paths between the on-premises data center and the AWS Cloud.
Which solution will meet these requirements while providing the HIGHEST throughput?

A. Configure a public VIF on the Direct Connect connectio
B. Configure an AWS Site-to-Site VPN connection to the transit gateway as a VPN attachment.
C. Configure a transit VIF on the Direct Connect connectio
D. Configure an IPsec VPN connection to an EC2 instance that is running third-party VPN software.
E. Configure MACsec for the Direct Connect connectio
F. Configure a transit VIF to a Direct Connect gateway that is associated with the transit gateway.
G. Configure a public VIF on the Direct Connect connectio
H. Configure two AWS Site-to-Site VPN connections to the transit gatewa
I. Enable equal-cost multi-path (ECMP) routing.

Answer: C

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/adding-macsec-security-to-aws-direct-connect-c

NEW QUESTION 15
A company uses a 1 Gbps AWS Direct Connect connection to connect its AWS environment to its
on-premises data center. The connection provides employees with access to an application VPC that is hosted on AWS. Many remote employees use a company-provided VPN to connect to the data center. These employees are reporting slowness when they access the application during business hours. On-premises users have started to report similar slowness while they are in the office.
The company plans to build an additional application on AWS. On-site and remote employees will use the additional application. After the deployment of this additional application, the company will need 20% more bandwidth than the company currently uses. With the increased usage, the company wants to add resiliency to the AWS connectivity. A network engineer must review the current implementation and must make improvements within a limited budget.
What should the network engineer do to meet these requirements MOST cost-effectively?

A. Set up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional traffic load from remote employees and the additional applicatio
B. Create a link aggregation group (LAG).
C. Deploy an AWS Site-to-Site VPN connection to the application VP
D. Configure the on-premises routing for the remote employees to connect to the Site-to-Site VPN connection.
E. Deploy Amazon Workspaces into the application VPInstruct the remote employees to connect to Workspaces.
F. Replace the existing 1 Gbps Direct Connect connection with two new 2 Gbps Direct Connect hosted connection
G. Create an AWS Client VPN endpoint in the application VP
H. Instruct the remote employees to connect to the Client VPN endpoint.

Answer: A

Explanation:
Setting up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional trafficload from remote employees and the additional application would provide more bandwidth and lower latency than a VPN connection over the public internet1. Creating a link aggregation group (LAG) with the existing and new Direct Connect connections would provide resiliency and redundancy for the AWS connectivity2.

NEW QUESTION 16
A global delivery company is modernizing its fleet management system. The company has several business units. Each business unit designs and maintains applications that are hosted in its own AWS account in separate application VPCs in the same AWS Region. Each business unit's applications are designed to get data from a central shared services VPC.
The company wants the network connectivity architecture to provide granular security controls. The architecture also must be able to scale as more business units consume data from the central shared services VPC in the future.
Which solution will meet these requirements in the MOST secure manner?

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

Answer: C

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

NEW QUESTION 18
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 19
A company is using Amazon Route 53 Resolver DNS Firewall in a VPC to block all domains except domains that are on an approved list. The company is concerned that if DNS Firewall is unresponsive, resources in the VPC might be affected if the network cannot resolve any DNS queries. To maintain application service level agreements, the company needs DNS queries to continue to resolve even if Route 53 Resolver does not receive a response from DNS Firewall.
Which change should a network engineer implement to meet these requirements?

A. Update the DNS Firewall VPC configuration to disable fail open for the VPC.
B. Update the DNS Firewall VPC configuration to enable fail open for the VPC.
C. Create a new DHCP options set with parameter dns_firewall_fail_open=fals
D. Associate the new DHCP options set with the VPC.
E. Create a new DHCP options set with parameter dns_firewall_fail_open=tru
F. Associate the new DHCP options set with the VPC.

Answer: B

NEW QUESTION 22
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 27
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 29
A company manages resources across VPCs in multiple AWS Regions. The company needs to connect to the resources by using its internal domain name. A network engineer needs to apply the aws.example.com DNS suffix to all resources.
What must the network engineer do to meet this requirement?

A. Create an Amazon Route 53 private hosted zone for aws.example.com in each Region that has resource
B. Associate the private hosted zone with that Region's VP
C. In the appropriate private hosted zone, create DNS records for the resources in each Region.
D. Create one Amazon Route 53 private hosted zone for aws.example.co
E. Configure the private hosted zone to allow zone transfers with every VPC.
F. Create one Amazon Route 53 private hosted zone for example.co
G. Create a single resource record for aws.example.com in the private hosted zon
H. Apply a multivalue answer routing policy to the recor
I. Add all VPC resources as separate values in the routing policy.
J. Create one Amazon Route 53 private hosted zone for aws.example.co
K. Associate the private hosted zone with every VPC that has resource
L. In the private hosted zone, create DNS records for all resources.

Answer: D

Explanation:
Creating one private hosted zone for aws.example.com and associating it with every VPC that has resources would enable DNS resolution for all resources by using their internal domain name. Creating an alias record in each private hosted zone with the full AWS service endpoint pointing to the interface VPC endpoint in the shared services VPC would enable private connectivity to Amazon S3 and AWS Systems Manager without using public endpoints.

NEW QUESTION 34
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 38
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 41
A company is planning to create a service that requires encryption in transit. The traffic must not be decrypted between the client and the backend of the service. The company will implement the service by using the gRPC protocol over TCP port 443. The service will scale up to thousands of simultaneous connections. The backend of the service will be hosted on an Amazon Elastic Kubernetes Service (Amazon EKS) duster with the Kubernetes Cluster Autoscaler and the Horizontal Pod Autoscaler configured. The company needs to use mutual TLS for two-way authentication between the client and the backend.
Which solution will meet these requirements?

A. Install the AWS Load Balancer Controller for Kubernete
B. Using that controller, configure a NetworkLoad Balancer with a TCP listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
C. Install the AWS Load Balancer Controller for Kubernete
D. Using that controller, configure an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
E. Create a target grou
F. Add the EKS managed node group's Auto Scaling group as a target Create an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the target group.
G. Create a target grou
H. Add the EKS managed node group’s Auto Scaling group as a targe
I. Create a Network Load Balancer with a TLS listener on port 443 to forward traffic to the target group.

Answer: B

Explanation:
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html#target-gro

NEW QUESTION 45
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 47
A company is deploying a new application in the AWS Cloud. The company wants a highly available web server that will sit behind an Elastic Load Balancer. The load balancer will route requests to multiple target groups based on the URL in the request. All traffic must use HTTPS. TLS processing must be offloaded to the load balancer. The web server must know the user’s IP address so that the company can keep accurate logs for security purposes.
Which solution will meet these requirements?

A. Deploy an Application Load Balancer with an HTTPS listene
B. Use path-based routing rules to forward the traffic to the correct target grou
C. Include the X-Forwarded-For request header with traffic to the targets.
D. Deploy an Application Load Balancer with an HTTPS listener for each domai
E. Use host-based routing rules to forward the traffic to the correct target group for each domai
F. Include the X-Forwarded-For request header with traffic to the targets.
G. Deploy a Network Load Balancer with a TLS listene
H. Use path-based routing rules to forward the traffic to the correct target grou
I. Configure client IP address preservation for traffic to the targets.
J. Deploy a Network Load Balancer with a TLS listener for each domai
K. Use host-based routing rules to forward the traffic to the correct target group for each domai
L. Configure client IP address preservation for traffic to the targets.

Answer: A

Explanation:
An Application Load Balancer (ALB) can be used to route traffic to multiple target groups based on the URL in the request. The ALB can be configured with an HTTPS listener to ensure all traffic uses HTTPS. TLS processing can be offloaded to the ALB, which reduces the load on the web server. Path-based routing rules can be used to route traffic to the correct target group based on the URL in the request. The X-Forwarded-For request header can be included with traffic to the targets, which will allow the web server to know the user's IP address and keep accurate logs for security purposes.

NEW QUESTION 52
A company's network engineer needs to design a new solution to help troubleshoot and detect network anomalies. The network engineer has configured Traffic Mirroring. However, the mirrored traffic is overwhelming the Amazon EC2 instance that is the traffic mirror target. The EC2 instance hosts tools that the company's security team uses to analyze the traffic. The network engineer needs to design a highly available solution that can scale to meet the demand of the mirrored traffic.
Which solution will meet these requirements?

A. Deploy a Network Load Balancer (NLB) as the traffic mirror targe
B. Behind the NL
C. deploy a fleet of EC2 instances in an Auto Scaling grou
D. Use Traffic Mirroring as necessary.
E. Deploy an Application Load Balancer (ALB) as the traffic mirror targe
F. Behind the ALB, deploy a fleet of EC2 instances in an Auto Scaling grou
G. Use Traffic Mirroring only during non-business hours.
H. Deploy a Gateway Load Balancer (GLB) as the traffic mirror targe
I. Behind the GL
J. deploy a fleet of EC2 instances in an Auto Scaling grou
K. Use Traffic Mirroring as necessary.
L. Deploy an Application Load Balancer (ALB) with an HTTPS listener as the traffic mirror targe
M. Behind the AL
N. deploy a fleet of EC2 instances in an Auto Scaling grou
O. Use Traffic Mirroring only during active events or business hours.

Answer: A

NEW QUESTION 53
A company deploys a new web application on Amazon EC2 instances. The application runs in private subnets in three Availability Zones behind an Application Load Balancer (ALB). Security auditors require encryption of all connections. The company uses Amazon Route 53 for DNS and uses AWS Certificate Manager (ACM) to automate SSL/TLS certificate provisioning. SSL/TLS connections are terminated on the ALB.
The company tests the application with a single EC2 instance and does not observe any problems. However, after production deployment, users report that they can log in but that they cannot use the application. Every new web request restarts the login process.
What should a network engineer do to resolve this issue?

A. Modify the ALB listener configuratio
B. Edit the rule that forwards traffic to the target grou
C. Change the rule to enable group-level stickines
D. Set the duration to the maximum application session length.
E. Replace the ALB with a Network Load Balance
F. Create a TLS listene
G. Create a new target group with the protocol type set to TLS Register the EC2 instance
H. Modify the target group configuration by enabling the stickiness attribute.
I. Modify the ALB target group configuration by enabling the stickiness attribut
J. Use an application-based cooki
K. Set the duration to the maximum application session length.
L. Remove the AL
M. Create an Amazon Route 53 rule with a failover routing policy for the application nam
N. Configure ACM to issue certificates for each EC2 instance.

Answer: C

NEW QUESTION 55
A Network Engineer is provisioning a subnet for a load balancer that will sit in front of a fleet of application servers in a private subnet. There is limited IP space left in the VPC CIDR. The application has few users now but is expected to grow quickly to millions of users.
What design will use the LEAST amount of IP space, while allowing for this growth?

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

Answer: C

NEW QUESTION 58
A company has deployed a new web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Amazon EC2 Auto Scaling group. Enterprise customers from around the world will use the application. Employees of these enterprise customers will connect to the application over HTTPS from office locations.
The company must configure firewalls to allow outbound traffic to only approved IP addresses. The employees of the enterprise customers must be able to access the application with the least amountof latency.
Which change should a network engineer make in the infrastructure to meet these requirements?

A. Create a new Network Load Balancer (NLB). Add the ALB as a target of the NLB.
B. Create a new Amazon CloudFront distributio
C. Set the ALB as the distribution's origin.
D. Create a new accelerator in AWS Global Accelerato
E. Add the ALB as an accelerator endpoint.
F. Create a new Amazon Route 53 hosted zon
G. Create a new record to route traffic to the ALB.

Answer: B

Explanation:
Amazon CloudFront is a content delivery network (CDN) that can speed up the delivery of static and dynamic web content, such as images, videos, and APIs2. CloudFront can also provide end-to-end encryption for HTTPS traffic by using SSL certificates from AWS Certificate Manager (ACM) or other sources2. CloudFron can also support session affinity (sticky sessions) with a load balancer-generated cookie or an application-based cookie policy2.

NEW QUESTION 60
A company uses a 4 Gbps AWS Direct Connect dedicated connection with a link aggregation group (LAG) bundle to connect to five VPCs that are deployed in the us-east-1 Region. Each VPC serves a different business unit and uses its own private VIF for connectivity to the on-premises environment. Users are reporting slowness when they access resources that are hosted on AWS.
A network engineer finds that there are sudden increases in throughput and that the Direct Connect connection becomes saturated at the same time for about an hour each business day. The company wants to know which business unit is causing the sudden increase in throughput. The network engineer must find out this information and implement a solution to resolve the problem.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 61
A company has deployed an application in a VPC that uses a NAT gateway for outbound traffic to the internet. A network engineer notices a large quantity of suspicious network traffic that is traveling from the VPC over the internet to IP addresses that are included on a deny list. The network engineer must implement a solution to determine which AWS resources are generating the suspicious traffic. The solution must minimize cost and administrative overhead.
Which solution will meet these requirements?

A. Launch an Amazon EC2 instance in the VP
B. Use Traffic Mirroring by specifying the NAT gateway as the source and the EC2 instance as the destinatio
C. Analyze the captured traffic by using open-source tools to identify the AWS resources that are generating the suspicious traffic.
D. Use VPC flow log
E. Launch a security information and event management (SIEM) solution in the VP
F. Configure the SIEM solution to ingest the VPC flow log
G. Run queries on the SIEM solution to identify the AWS resources that are generating the suspicious traffic.
H. Use VPC flow log
I. Publish the flow logs to a log group in Amazon CloudWatch Log
J. Use CloudWatch Logs Insights to query the flow logs to identify the AWS resources that are generating the suspicious traffic.
K. Configure the VPC to stream the network traffic directly to an Amazon Kinesis data strea
L. Send the data from the Kinesis data stream to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Athena to query the data to identify the AWS resources that are generating the suspicious traffic.

Answer: C

NEW QUESTION 65
A company wants to improve visibility into its AWS environment. The AWS environment consists of multiple VPCs that are connected to a transit gateway. The transit gateway connects to an on-premises data center through an AWS Direct Connect gateway and a pair of redundant Direct Connect connections that use transit VIFs. The company must receive notification each time a new route is advertised to AWS from on premises over Direct Connect.
What should a network engineer do to meet these requirements?

A. Enable Amazon CloudWatch metrics on Direct Connect to track the received route
B. Configure a CloudWatch alarm to send notifications when routes change.
C. Onboard Transit Gateway Network Manager to Amazon CloudWatch Logs Insight
D. Use Amazon EventBridge (Amazon CloudWatch Events) to send notifications when routes change.
E. Configure an AWS Lambda function to periodically check the routes on the Direct Connect gateway and to send notifications when routes change.
F. Enable Amazon CloudWatch Logs on the transit VIFs to track the received route

G. Create a metric filter Set an alarm on the filter to send notifications when routes change.

Answer: B

Explanation:
https://docs.aws.amazon.com/network-manager/latest/cloudwan/cloudwan-cloudwatch-events.html
To receive notification each time a new route is advertised to AWS from on premises over Direct Connect, a network engineer should onboard Transit Gateway Network Manager to Amazon CloudWatch Logs Insights and use Amazon EventBridge (Amazon CloudWatch Events) to send notifications when routes change (Option B). This solution allows for real-time monitoring of route changes and automatic notification when new routes are advertised.

NEW QUESTION 67
A company has multiple AWS accounts. Each account contains one or more VPCs. A new security guideline requires the inspection of all traffic between VPCs. The company has deployed a transit gateway that provides connectivity between all VPCs. The company also has deployed a shared services VPC with Amazon EC2 instances that include IDS services for stateful inspection. The EC2 instances are deployed across three Availability Zones. The company has set up VPC associations and routing on the transit gateway. The company has migrated a few test VPCs to the new solution for traffic inspection.
Soon after the configuration of routing, the company receives reports of intermittent connections for traffic that crosses Availability Zones.
What should a network engineer do to resolve this issue?

A. Modify the transit gateway VPC attachment on the shared services VPC by enabling cross-Availability Zone load balancing.
B. Modify the transit gateway VPC attachment on the shared services VPC by enabling appliance mode support.
C. Modify the transit gateway by selecting VPN equal-cost multi-path (ECMP) routing support.
D. Modify the transit gateway by selecting multicast support.

Answer: B

Explanation:
To resolve the issue of intermittent connections for traffic that crosses Availability Zonesafter configuring routing for traffic inspection between VPCs using a transit gateway and EC2 instances with IDS services in a shared services VPC, a network engineer should modify the transit gateway VPC attachment on the shared services VPC by enabling appliance mode support (Option B). This will ensure that traffic is routed to the same EC2 instance for stateful inspection and prevent intermittent connections.

NEW QUESTION 69
A security team is performing an audit of a company's AWS deployment. The security team is concerned that two applications might be accessing resources that should be blocked by network ACLs and security groups. The applications are deployed across two Amazon Elastic Kubernetes Service (Amazon EKS) clusters that use the Amazon VPC Container Network Interface (CNI) plugin for Kubernetes. The clusters are in separate subnets within the same VPC and have a Cluster Autoscaler configured.
The security team needs to determine which POD IP addresses are communicating with which services throughout the VPC. The security team wants to limit the number of flow logs and wants to examine the traffic from only the two applications.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create VPC flow logs in the default forma
B. Create a filter to gather flow logs only from the EKS nodes.Include the srcaddr field and the dstaddr field in the flow logs.
C. Create VPC flow logs in a custom forma
D. Set the EKS nodes as the resource Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.
E. Create VPC flow logs in a custom forma
F. Set the application subnets as resource
G. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.
H. Create VPC flow logs in a custom forma
I. Create a filter to gather flow logs only from the EKS nodes.Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.

Answer: D

NEW QUESTION 74
A company delivers applications over the internet. An Amazon Route 53 public hosted zone is the
authoritative DNS service for the company and its internet applications, all of which are offered from the same domain name.
A network engineer is working on a new version of one of the applications. All the application's components are hosted in the AWS Cloud. The application has a three-tier design. The front end is delivered through Amazon EC2 instances that are deployed in public subnets with Elastic IP addresses assigned. The backend components are deployed in private subnets from RFC1918.
Components of the application need to be able to access other components of the application within the application's VPC by using the same host names as the host names that are used over the public internet. The network engineer also needs to accommodate future DNS changes, such as the introduction of new host names or the retirement of DNS entries.
Which combination of steps will meet these requirements? (Choose three.)

A. Add a geoproximity routing policy in Route 53.
B. Create a Route 53 private hosted zone for the same domain name Associate the application’s VPC with the new private hosted zone.
C. Enable DNS hostnames for the application's VPC.
D. Create entries in the private hosted zone for each name in the public hosted zone by using the corresponding private IP addresses.
E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs when AWS CloudTrail logs a Route 53 API call to the public hosted zon
F. Create an AWS Lambda function as the target of the rul
G. Configure the function to use the event information to update the privatehosted zone.
H. Add the private IP addresses in the existing Route 53 public hosted zone.

Answer: BCD

NEW QUESTION 75
A company is hosting an application on Amazon EC2 instances behind a Network Load Balancer (NLB). A solutions architect added EC2 instances in a second Availability Zone to improve the availability of the application. The solutions architect added the instances to the NLB target group.
The company's operations team notices that traffic is being routed only to the instances in the first Availability Zone.
What is the MOST operationally efficient solution to resolve this issue?

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

Answer: A

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

NEW QUESTION 76
An Australian ecommerce company hosts all of its services in the AWS Cloud and wants to expand its customer base to the United States (US). The company is targeting the western US for the expansion.
The company's existing AWS architecture consists of four AWS accounts with multiple VPCs deployed in the ap-southeast-2 Region. All VPCs are attached to a transit gateway in ap-southeast-2. There are dedicated VPCs for each application service. The company also has VPCs for centralized security features such as proxies, firewalls, and logging.
The company plans to duplicate the infrastructure from ap-southeast-2 to the us-west-1 Region. A network engineer must establish connectivity between the various applications in the two Regions. The solution must maximize bandwidth, minimize latency and minimize operational overhead.
Which solution will meet these requirements?

A. Create VPN attachments between the two transit gateway
B. Configure the VPN attachments to use BGP routing between the two transit gateways.
C. Peer the transit gateways in each Regio
D. Configure routing between the two transit gateways for each Region's IP addresses.
E. Create a VPN server in a VPC in each Regio
F. Update the routing to point to the VPN servers for the IP addresses in alternate Regions.
G. Attach the VPCs in us-west-1 to the transit gateway in ap-southeast-2.

Answer: B

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

NEW QUESTION 77
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 82
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 86
Your security team implements a host-based firewall on all of your Amazon Elastic Compute Cloud (EC2) instances to block all outgoing traffic. Exceptions must be requested for each specific requirement. Until you request a new rule, you cannot access the instance metadata service. Which firewall rule should you request to be added to your instances to allow instance metadata access?

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

Answer: C

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

NEW QUESTION 91
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 96
A company is migrating an existing application to a new AWS account. The company will deploy the application in a single AWS Region by using one VPC and

multiple Availability Zones. The application will run on Amazon EC2 instances. Each Availability Zone will have several EC2 instances. The EC2 instances will be deployed in private subnets.
The company's clients will connect to the application by using a web browser with the HTTPS protocol. Inbound connections must be distributed across the Availability Zones and EC2 instances. All connections from the same client session must be connected to the same EC2 instance. The company must provide end-to-end encryption for all connections between the clients and the application by using the application SSL certificate.
Which solution will meet these requirements?

A. Create a Network Load Balance
B. Create a target grou
C. Set the protocol to TCP and the port to 443 for the target grou
D. Turn on session affinity (sticky sessions). Register the EC2 instances as target
E. Create a listene
F. Set the protocol to TCP and the port to 443 for the listene
G. Deploy SSL certificates to the EC2 instances.
H. Create an Application Load Balance
I. Create a target grou
J. Set the protocol to HTTP and the port to 80 for the target grou
K. Turn on session affinity (sticky sessions) with an application-based cookie polic
L. Register the EC2 instances as target
M. Create an HTTPS listene
N. Set the default action to forward to the target grou
O. Use AWS Certificate Manager (ACM) to create a certificatefor the listener.
P. Create a Network Load Balance
Q. Create a target grou
R. Set the protocol to TLS and the port to 443 for the target grou
S. Turn on session affinity (sticky sessions). Register the EC2 instances as target
T. Create a listene
. Set the protocol to TLS and the port to 443 for the listene
. Use AWS Certificate Manager (ACM) to create a certificate for the application.
. Create an Application Load Balance
. Create a target grou
. Set the protocol to HTTPS and the port to 443 for the target grou
. Turn on session affinity (sticky sessions) with an application-based cookie polic
. Register the EC2 instances as target
. Create an HTTP listene
. Set the port to 443 for the listene
. Set the default action to forward to the target group.

Answer: A

NEW QUESTION 97
.....

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