

Amazon

Exam Questions AWS-Certified-Advanced-Networking-Specialty

Amazon AWS Certified Advanced Networking - Specialty



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 interface
- 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 interface
- H. Send the traffic to an additional EC2 instance
- 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 interface
- K. Publish the logs to an Amazon S3 bucket
- L. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure
- M. Use Athena to query and analyze the logs.
- N. Enable NAT gateway access log
- O. Publish the logs to an Amazon S3 bucket
- P. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure
- 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 company is running multiple workloads on Amazon EC2 instances in public subnets. In a recent incident, an attacker exploited an application vulnerability on one of the EC2 instances to gain access to the instance. The company fixed the application and launched a replacement EC2 instance that contains the updated application.

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 3

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 VPC
- B. Connect the inspection VPC to the transit gateway by using a VPC attachment
- C. Create a target group, and register the appliances with the target group
- D. Create a Network Load Balancer (NLB), and set it up to forward to the newly created target group
- 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 VPC
- G. Connect the inspection VPC to the transit gateway by using a VPC attachment
- H. Create a target group, and register the appliances with the target group
- I. Create a Gateway Load Balancer, and set it up to forward to the newly created target group

- 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 gateway
- L. Associate one route table with all the attachments of the application VPC
- M. Associate the other route table with the inspection VPC's attachments
- N. Propagate all VPC attachments into the inspection route table
- O. Define a static default route in the application route table
- P. Enable appliance mode on the attachment that connects the inspection VPC.
- Q. Configure two route tables on the transit gateway
- R. Associate one route table with all the attachments of the application VPC
- S. Associate the other route table with the inspection VPC's attachments
- T. Propagate all VPC attachments into the application route table
- . Define a static default route in the inspection route table
- . Enable appliance mode on the attachment that connects the inspection VPC.
- . Configure one route table on the transit gateway
- . Associate the route table with all the VPC
- . Propagate all VPC attachments into the route table
- . Define a static default route in the route table.

Answer: BC

NEW QUESTION 4

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

Answer: A

NEW QUESTION 5

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 6

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

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

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

- A. Create a VPN connection over the Direct Connect connection by using the on-premises firewall
- 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 deployed
- G. Specify the NAT gateway type as private
- 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 7

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 VPC
- C. Reboot the Amazon Linux 2 EC2 instance.
- D. Create an Amazon Route 53 Resolver rule
- E. Associate the rule with the VPC
- 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 hosts file in the Amazon Linux 2 EC2 instance in the VPC to map 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 VPC
- 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). 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 endpoints.

NEW QUESTION 8

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 this centralized 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 KMS
- B. Modify the interface endpoint by disabling the private DNS name
- C. Create a private hosted zone in the shared services account with an alias record that points to the interface endpoint
- 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 KMS
- F. Modify the interface endpoint by disabling the private DNS name
- G. Create a private hosted zone in each spoke AWS account with an alias record that points to the interface endpoint
- H. Associate each private hosted zone with the shared services AWS account.
- I. In each spoke AWS account, create an interface endpoint for AWS KMS
- J. Modify each interface endpoint by disabling the private DNS name
- K. Create a private hosted zone in each spoke AWS account with an alias record that points to each interface endpoint
- L. Associate each private hosted zone with the shared services AWS account.
- M. In each spoke AWS account, create an interface endpoint for AWS KMS

- 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 endpoint
- P. Associate the private hosted zone with the spoke VPCs in each AWS account.

Answer: A

NEW QUESTION 9

A network engineer must develop an AWS CloudFormation template that can create a virtual private gateway, a customer gateway, a VPN connection, and static routes in a route table. During testing of the template, the network engineer notes that the CloudFormation template has encountered an error and is rolling back. What should the network engineer do to resolve the error?

- A. Change the order of resource creation in the CloudFormation template.
- B. Add the DependsOn attribute to the resource declaration for the virtual private gateway
- C. Specify the route table entry resource.
- D. Add a wait condition in the template to wait for the creation of the virtual private gateway.
- E. Add the DependsOn attribute to the resource declaration for the route table entry
- F. Specify the virtual private gateway resource.

Answer: D

NEW QUESTION 10

A company is planning to use Amazon S3 to archive financial data. The data is currently stored in an on-premises data center. The company uses AWS Direct Connect with a Direct Connect gateway and a transit gateway to connect to the on-premises data center. The data cannot be transported over the public internet and must be encrypted in transit. Which solution will meet these requirements?

- A. Create a Direct Connect public VIF
- B. Set up an IPsec VPN connection over the public VIF to access Amazon S3. Use HTTPS for communication.
- C. Create an IPsec VPN connection over the transit VIF
- D. Create a VPC and attach the VPC to the transit gateway
- E. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.
- F. Create a VPC and attach the VPC to the transit gateway
- G. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.
- H. Create a Direct Connect public VIF
- I. Set up an IPsec VPN connection over the public VIF to the transit gateway
- J. Create an attachment for Amazon S3. Use HTTPS for communication.

Answer: B

Explanation:

<https://docs.aws.amazon.com/vpn/latest/s2svpn/private-ip-dx.html>

An IPsec VPN connection over the transit VIF can encrypt traffic between the on-premises network and AWS without using public IP addresses or the internet. A VPC endpoint for Amazon S3 can enable private access to S3 buckets within the same region. HTTPS can provide additional encryption for communication.

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 VIF
- B. Associate the transit gateway in us-east-1 with the same Direct Connect gateway
- C. Enable SiteLink for the transit VIF and the private VIF.
- D. Connect the VPC in eu-west-2 to a new transit gateway
- E. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VIF
- F. Associate the transit gateway in us-east-1 with the same Direct Connect gateway
- 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 gateway
- J. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VIF
- K. Create a new Direct Connect gateway
- L. Associate the transit gateway in us-east-1 with the new Direct Connect gateway
- 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 VIF
- P. Create a new Direct Connect gateway
- Q. Associate the transit gateway in us-east-1 with the new Direct Connect gateway
- R. Enable SiteLink for the transit VIF and the private VIF.

Answer: C

NEW QUESTION 13

A company has deployed a web application on AWS. The web application uses an Application Load Balancer (ALB) across multiple Availability Zones. The targets of the ALB are AWS Lambda functions. The web application also uses Amazon CloudWatch metrics for monitoring. Users report that parts of the web application are not loading properly. A network engineer needs to troubleshoot the problem. The network engineer enables access logging for the ALB. What should the network engineer do next to determine which errors the ALB is receiving?

- A. Send the logs to Amazon CloudWatch Log
- B. Review the ALB logs in CloudWatch Insights to determine which error messages the ALB is receiving.
- C. Configure the Amazon S3 bucket destination
- D. Use Amazon Athena to determine which error messages the ALB is receiving.
- E. Configure the Amazon S3 bucket destination
- F. After Amazon CloudWatch Logs pulls the ALB logs from the S3 bucket automatically, review the logs in CloudWatch Logs to determine which error messages the ALB is receiving.
- G. Send the logs to Amazon CloudWatch Log
- H. Use the Amazon Athena CloudWatch Connector to determine which error messages the ALB is receiving.

Answer: B

Explanation:

Access logs is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logs for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify as compressed files. You can disable access logs at any time. <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

NEW QUESTION 15

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 18

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 21

A company has two AWS accounts one for Production and one for Connectivity. A network engineer needs to connect the Production account VPC to a transit gateway in the Connectivity account. The feature to auto accept shared attachments is not enabled on the transit gateway. Which set of steps should the network engineer follow in each AWS account to meet these requirements?

- A. * 1. In the Production account: Create a resource share in AWS Resource Access Manager for the transit gateway
- B. Provide the Connectivity account ID
- 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 attachment
- 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 ID
- 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 attachment
- 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 ID
- 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 attachment
- 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 gateway
- 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 attachment
- 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 23

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 has monitoring 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 than $0.02 * OUT$, 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 address
- B. Create a new VPN connection that supports IPv6 connectivity
- C. Add an egress-only internet gateway
- 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 address
- F. Update the existing VPN connection to support IPv6 connectivity
- G. Add an egress-only internet gateway
- 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 address
- J. Create a new VPN connection that supports IPv6 connectivity
- K. Add an egress-only internet gateway
- 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 address
- N. Create a new VPN connection that supports IPv6 connectivity
- O. Add a NAT gateway
- 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 30

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 application
- B. Create a link aggregation group (LAG).
- C. Deploy an AWS Site-to-Site VPN connection to the application VPC
- 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 VPC and instruct the remote employees to connect to Workspaces.
- F. Replace the existing 1 Gbps Direct Connect connection with two new 2 Gbps Direct Connect hosted connections
- G. Create an AWS Client VPN endpoint in the application VPC
- 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 traffic load from remote employees and the additional application would provide more bandwidth and lower latency than a VPN connection over the public internet. Creating a link aggregation group (LAG) with the existing and new Direct Connect connections would provide resiliency and redundancy for the AWS connectivity.

NEW QUESTION 31

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 NLB. Use 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 ALB
- G. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator
- H. Place the EC2 instances behind an Amazon CloudFront distribution
- I. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

Answer: B

NEW QUESTION 33

A company has expanded its network to the AWS Cloud by using a hybrid architecture with multiple AWS accounts. The company has set up a shared AWS account for the connection to its on-premises data centers and the company offices. The workloads consist of private web-based services for internal use. These services run in different AWS accounts. Office-based employees consume these services by using a DNS name in an on-premises DNS zone that is named example.internal.

The process to register a new service that runs on AWS requires a manual and complicated change request to the internal DNS. The process involves many teams.

The company wants to update the DNS registration process by giving the service creators access that will allow them to register their DNS records. A network engineer must design a solution that will achieve this goal. The solution must maximize cost-effectiveness and must require the least possible number of configuration changes.

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

- A. Create a record for each service in its local private hosted zone (serviceA.account1.aws.example.internal). Provide this DNS record to the employees who need access.
- B. Create an Amazon Route 53 Resolver inbound endpoint in the shared account VPC
- C. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS server
- D. Set the forwarding IP addresses to the inbound endpoint's IP addresses that were created.
- E. Create an Amazon Route 53 Resolver rule to forward any queries made to onprem.example.internal to the on-premises DNS servers.
- F. Create an Amazon Route 53 private hosted zone named aws.example.internal in the shared AWS account to resolve queries for this domain.
- G. Launch two Amazon EC2 instances in the shared AWS account
- H. Install BIND on each instance
- I. Create a DNS conditional forwarder on each BIND server to forward queries for each subdomain under aws.example.internal to the appropriate private hosted zone in each AWS account
- J. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS server
- K. Set the forwarding IP addresses to the IP addresses of the BIND servers.
- L. Create a private hosted zone in the shared AWS account for each account that runs the service. Configure the private hosted zone to contain aws.example.internal in the domain (account1.aws.example.internal). Associate the private hosted zone with the VPC that runs the service and the shared account VPC.

Answer: ABD

Explanation:

To meet the requirements of updating the DNS registration process while maximizing cost-effectiveness and minimizing configuration changes, the network engineer should take the following steps:

- Create an Amazon Route 53 Resolver inbound endpoint in the shared account VPC. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS servers. Set the forwarding IP addresses to the inbound endpoint's IP addresses that were created (Option B).
- Create an Amazon Route 53 private hosted zone named aws.example.internal in the shared AWS account to resolve queries for this domain (Option D).
- Create a record for each service in its local private hosted zone (serviceA.account1.aws.example.internal). Provide this DNS record to the employees who need access (Option A).

These steps will allow service creators to register their DNS records while keeping costs low and minimizing configuration changes.

NEW QUESTION 34

A global delivery company is modernizing its fleet management system. The company has several business units. Each business unit designs and maintains applications that are hosted in its own AWS account in separate application VPCs in the same AWS Region. Each business unit's applications are designed to get data from a central shared services VPC.

The company wants the network connectivity architecture to provide granular security controls. The architecture also must be able to scale as more business units consume data from the central shared services VPC in the future.

Which solution will meet these requirements in the MOST secure manner?

- A. Create a central transit gateway
- B. Create a VPC attachment to each application VPC
- 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 VPC. Create VPC endpoints in each application VPC.
- F. Create a central transit VPC with a VPN appliance from AWS Marketplace
- G. Create a VPN attachment from each VPC to the transit VPC
- 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 37

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 appliance
- 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 42

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 multivalued 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 46

A company is using a NAT gateway to allow internet connectivity for private subnets in a VPC in the us-west-2 Region. After a security audit, the company needs to remove the NAT gateway.

In the private subnets, the company has resources that use the unified Amazon CloudWatch agent. A network engineer must create a solution to ensure that the unified CloudWatch agent continues to work after the removal of the NAT gateway.

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

- A. Validate that private DNS is enabled on the VPC by setting the enableDnsHostnames VPC attribute and the enableDnsSupport VPC attribute to true.
- B. Create a new security group with an entry to allow outbound traffic that uses the TCP protocol on port 443 to destination 0.0.0.0/0
- C. Create a new security group with entries to allow inbound traffic that uses the TCP protocol on port 443 from the IP prefixes of the private subnets.
- D. Create the following interface VPC endpoints in the VPC: com.amazonaws.us-west-2.logs and com.amazonaws.us-west-2.monitorin
- E. Associate the new security group with the endpoint network interfaces.
- F. Create the following interface VPC endpoint in the VPC: com.amazonaws.us-west-2.cloudwatch. Associate the new security group with the endpoint network interfaces.
- G. Associate the VPC endpoint or endpoints with route tables that the private subnets use.

Answer: BDF

NEW QUESTION 47

A company is hosting an application on Amazon EC2 instances behind an Application Load Balancer. The instances are in an Amazon EC2 Auto Scaling group. Because of a recent change to a security group, external users cannot access the application.

A network engineer needs to prevent this downtime from happening again. The network engineer must implement a solution that remediates noncompliant changes to security groups.

Which solution will meet these requirements?

- A. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuratio
- B. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.
- C. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuratio
- D. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.
- E. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuratio
- F. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.
- G. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuratio
- H. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.

Answer: D

Explanation:

Configuring an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration would enable evaluation of the compliance status of the security groups based on predefined or custom rules³. Creating an AWS Systems Manager Automation runbook to remediate noncompliant security groups would enable automation of the remediation process². Additionally, configuring AWS Config to trigger the runbook when a noncompliant change is detected would enable timely and consistent remediation of security group changes.

NEW QUESTION 51

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

Which solution will meet these requirements?

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

Answer: B

NEW QUESTION 54

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

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

What should the network engineer do to meet these requirements?

- A. Remove the Direct Connect gateway, and create a new private virtual interface from each company router to the virtual private gateway of the VPC.
- B. Change the router configurations to summarize the advertised routes.
- C. Open a support ticket to increase the quota on advertised routes to the VPC route table.
- D. Create an AWS Transit Gateway
- 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 57

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 table
- C. Propagate all application VPC attachments to this transit gateway route table.
- D. Configure a separate transit gateway route table for each application VPC
- E. Associate each application VPC attachment with its respective transit gateway route table
- 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 VPCs
- H. Associate all application VPCs with this transit gateway route table
- I. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.
- J. Configure a separate transit gateway route table for the shared services VPC
- K. Associate the shared services VPC attachment with this transit gateway route table
- 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 VPC
- N. Associate the VPN attachment and the shared services VPC attachment with this transit gateway route table
- O. Propagate all application VPC attachments to this transit gateway route table.

Answer: BD

NEW QUESTION 59

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) cluster with the Kubernetes Cluster Autoscaler and the Horizontal Pod Autoscaler configured. The company needs to use mutual TLS for two-way authentication between the client and the backend. Which solution will meet these requirements?

- A. Install the AWS Load Balancer Controller for Kubernetes
- B. Using that controller, configure a Network Load Balancer with a TCP listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
- C. Install the AWS Load Balancer Controller for Kubernetes
- 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 group
- 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 group
- H. Add the EKS managed node group's Auto Scaling group as a target
- 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-group>

NEW QUESTION 60

A company has a global network and is using transit gateways to connect AWS Regions together. The company finds that two Amazon EC2 instances in different Regions are unable to communicate with each other. A network engineer needs to troubleshoot this connectivity issue. What should the network engineer do to meet this requirement?

- A. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables and in the VPC route table
- B. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.
- C. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables. Verify that the VPC route tables are correct
- D. Use AWS Firewall Manager to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.
- E. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables. Verify that the VPC route tables are correct
- F. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.
- G. Use VPC Reachability Analyzer to analyze routes in the transit gateway route table
- H. Verify that the VPC route tables are correct
- I. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.

Answer: C

Explanation:

Using AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between VPCs and transit gateways¹. Verifying that the VPC route tables are correct would enable identification of routing issues within a VPC. Using VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC would enable identification of traffic filtering issues within a VPC². Additionally, using VPC Reachability Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between transit gateways in different Regions. VPC Reachability Analyzer is a configuration analysis tool that enables connectivity testing between a source resource and a destination resource in your VPCs.

NEW QUESTION 65

A software company offers a software-as-a-service (SaaS) accounting application that is hosted in the AWS Cloud. The application requires connectivity to the company's on-premises network. The company has two redundant 10 GB AWS Direct Connect connections between AWS and its on-premises network to accommodate the growing demand for the application.

The company already has encryption between its on-premises network and the colocation. The company needs to encrypt traffic between AWS and the edge routers in the colocation within the next few months. The company must maintain its current bandwidth.

What should a network engineer do to meet these requirements with the LEAST operational overhead?

- A. Deploy a new public VIF with encryption on the existing Direct Connect connection
- B. Reroute traffic through the new public VIF.
- C. Create a virtual private gateway. Deploy new AWS Site-to-Site VPN connections from on-premises to the virtual private gateway. Reroute traffic from the Direct Connect private VIF to the new VPNs.
- D. Deploy a new pair of 10 GB Direct Connect connections with MACsec
- E. Configure MACsec on the edge router
- F. Reroute traffic to the new Direct Connect connection
- G. Decommission the original Direct Connect connections
- H. Deploy a new pair of 10 GB Direct Connect connections with MACsec
- I. Deploy a new public VIF on the new Direct Connect connection
- J. Deploy two AWS Site-to-Site VPN connections on top of the new public VIF
- K. Reroute traffic from the existing private VIF to the new Site-to-Site connection
- L. Decommission the original Direct Connect connections.

Answer: C

NEW QUESTION 67

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 69

An organization launched an IPv6-only web portal to support IPv6-native mobile clients. Front-end instances launch in an Amazon VPC associated with an appropriate IPv6 CIDR. The VPC IPv4 CIDR is fully utilized. A single subnet exists in each of two Availability Zones with appropriately configured IPv6 CIDR associations. Auto Scaling is properly configured, and no Elastic Load Balancing is used.

Customers say the service is unavailable during peak load times. The network engineer attempts to launch an instance manually and receives the following message: "There are not enough free addresses in subnet 'subnet-12345677' to satisfy the requested number of instances."

What action will resolve the availability problem?

- A. Create a new subnet using a VPC secondary IPv6 CIDR, and associate an IPv6 CID
- B. Include the new subnet in the Auto Scaling group.
- C. Create a new subnet using a VPC secondary IPv4 CIDR, and associate an IPv6 CID
- D. Include the new subnet in the Auto Scaling group.
- E. Resize the IPv6 CIDR on each of the existing subnet
- F. Modify the Auto Scaling group maximum number of instances.
- G. Add a secondary IPv4 CIDR to the Amazon VP
- H. Assign secondary IPv4 address space to each of the existing subnets.

Answer: B

NEW QUESTION 71

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 VP Add a route to the existing subnet route tables to point 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 traffic
- H. Associate the security group with the egress-only internet gateway.

Answer: C

NEW QUESTION 73

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

Which solution will meet these requirements?

- A. Create a new VPC for the SD-WAN hub virtual appliance
- B. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway
- C. Configure BGP over the IPsec VPN connections
- D. Assign a new CIDR block to the transit gateway
- E. Create a new VPC for the SD-WAN hub virtual appliance
- F. Attach the new VPC to the transit gateway with a VPC attachment
- G. Add a transit gateway Connect attachment
- H. Create a Connect peer and specify the GRE and BGP parameter
- I. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.
- J. Create a new VPC for the SD-WAN hub virtual appliance
- K. Attach the new VPC to the transit gateway with a VPC attachment
- L. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway
- M. Configure BGP over the IPsec VPN connections.
- N. Assign a new CIDR block to the transit gateway
- O. Create a new VPC for the SD-WAN hub virtual appliance
- P. Attach the new VPC to the transit gateway with a VPC attachment
- Q. Add a transit gateway Connect attachment
- R. Create a Connect peer and specify the VXLAN and BGP parameter
- S. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.

Answer: D

NEW QUESTION 77

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 82

An international company provides early warning about tsunamis. The company plans to use IoT devices to monitor sea waves around the world. The data that is collected by the IoT devices must reach the company's infrastructure on AWS as quickly as possible. The company is using three operation centers around the world. Each operation center is connected to AWS through its own AWS Direct Connect connection. Each operation center is connected to the internet through at least two upstream internet service providers.

The company has its own provider-independent (PI) address space. The IoT devices use TCP protocols for reliable transmission of the data they collect. The IoT devices have both landline and mobile internet connectivity. The infrastructure and the solution will be deployed in multiple AWS Regions. The company will use Amazon Route 53 for DNS services.

A network engineer needs to design connectivity between the IoT devices and the services that run in the AWS Cloud.

Which solution will meet these requirements with the HIGHEST availability?

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

Answer: B

Explanation:

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

NEW QUESTION 86

A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin in an Amazon CloudFront distribution. The company wants to implement a custom authentication system that will provide a token for its authenticated customers.

The web application must ensure that the GET/POST requests come from authenticated customers before it delivers the content. A network engineer must design a solution that gives the web application the ability to identify authorized customers.

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

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

Answer: C

NEW QUESTION 89

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 Config
- 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 Inventory
- 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 compliance.
 Creating rules that reflect the desired configuration settings would enable evaluation of whether the network resources comply with network security policies.
 Setting remediation for noncompliant resources would enable automatic correction of undesired configurations.

NEW QUESTION 90

A company operates its IT services through a multi-site hybrid infrastructure. The company deploys resources on AWS in the us-east-1 Region and in the eu-west-2 Region. The company also deploys resources in its own data centers that are located in the United States (US) and in the United Kingdom (UK). In both AWS Regions, the company uses a transit gateway to connect 15 VPCs to each other. The company has created a transit gateway peering connection between the two transit gateways. The VPC CIDR blocks do not overlap with each other or with IP addresses used within the data centers. The VPC CIDR prefixes can also be aggregated either on a Regional level or for the company's entire AWS environment.

The data centers are connected to each other by a private WAN connection. IP routing information is exchanged dynamically through Interior BGP (iBGP) sessions. The data centers maintain connectivity to AWS through one AWS Direct Connect connection in the US and one Direct Connect connection in the UK. Each Direct Connect connection is terminated on a Direct Connect gateway and is associated with a local transit gateway through a transit VIF.

Traffic follows the shortest geographical path from source to destination. For example, packets from the UK data center that are targeted to resources in eu-west-2 travel across the local Direct Connect connection. In cases of cross-Region data transfers, such as from the UK data center to VPCs in us-east-1, the private WAN connection must be used to minimize costs on AWS. A network engineer has configured each transit gateway association on the Direct Connect gateway to advertise VPC-specific CIDR IP prefixes only from the local Region. The routes toward the other Region must be learned through BGP from the routers in the other data center in the original, non-aggregated form.

The company recently experienced a problem with cross-Region data transfers because of issues with its private WAN connection. The network engineer needs to modify the routing setup to prevent similar interruptions in the future. The solution cannot modify the original traffic routing goal when the network is operating normally.

Which modifications will meet these requirements? (Choose two.)

- A. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connection
- B. Add the company's entire AWS environment aggregate route to the list of subnets advertised through the local Direct Connect connection.
- C. Add the CIDR prefixes from the other Region VPCs and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection
- D. Configure data center routers to make routing decisions based on the BGP communities received.
- E. Add the aggregate IP prefix for the other Region and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection.
- F. Add the aggregate IP prefix for the company's entire AWS environment and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection.
- G. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connection
- H. Add both Regional aggregate IP prefixes to the list of subnets advertised through the Direct Connect connection on both sides of the network
- I. Configure data center routers to make routing decisions based on the BGP communities received.

Answer: AD

NEW QUESTION 94

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

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

What should a network engineer do to resolve this issue?

- A. Modify the ALB listener configuration
- B. Edit the rule that forwards traffic to the target group
- C. Change the rule to enable group-level stickiness
- D. Set the duration to the maximum application session length.
- E. Replace the ALB with a Network Load Balance
- F. Create a TLS listener
- 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 attribute
- J. Use an application-based cookie
- K. Set the duration to the maximum application session length.
- L. Remove the ALB
- M. Create an Amazon Route 53 rule with a failover routing policy for the application name
- N. Configure ACM to issue certificates for each EC2 instance.

Answer: C

NEW QUESTION 99

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

How should the network engineer set up the Direct Connect connection to meet these requirements?

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

H. Create one dedicated connectio

I. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use two Direct Connect gateways, one for each VIF, to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.

Answer: B

Explanation:

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

NEW QUESTION 104

An organization is replacing a tape backup system with a storage gateway. there is currently no connectivity to AWS. Initial testing is needed. What connection option should the organization use to get up and running at minimal cost?

- A. Use an internet connection.
- B. Set up an AWS VPN connection.
- C. Provision an AWS Direct Connection private virtual interface.
- D. Provision a Direct Connect public virtual interface.

Answer: A

NEW QUESTION 108

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 110

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 rules. 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 ranges. Updating security groups to use the prefix list instead of the partner IP address range would enable simplification of security group rules. 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 list.

NEW QUESTION 115

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 format.
- 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 format.
- 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 format.
- 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 format.
- 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 119

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 zone.
- F. Create an AWS Lambda function as the target of the rule.
- G. Configure the function to use the event information to update the private hosted zone.
- H. Add the private IP addresses in the existing Route 53 public hosted zone.

Answer: BCD

NEW QUESTION 121

A company recently migrated its Amazon EC2 instances to VPC private subnets to satisfy a security compliance requirement. The EC2 instances now use a NAT gateway for internet access. After the migration, some long-running database queries from private EC2 instances to a publicly accessible third-party database no longer receive responses. The database query logs reveal that the queries successfully completed after 7 minutes but that the client EC2 instances never received the response.

Which configuration change should a network engineer implement to resolve this issue?

- A. Configure the NAT gateway timeout to allow connections for up to 600 seconds.
- B. Enable enhanced networking on the client EC2 instances.
- C. Enable TCP keepalive on the client EC2 instances with a value of less than 300 seconds.
- D. Close idle TCP connections through the NAT gateway.

Answer: C

Explanation:

When a TCP connection is idle for a long time, it may be terminated by network devices, including the NAT gateway. By enabling TCP keepalive, the client EC2 instances can periodically send packets to the third-party database to indicate that the connection is still active, preventing it from being terminated prematurely.

NEW QUESTION 126

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

The company's existing AWS architecture consists of four AWS accounts with multiple VPCs deployed in the ap-southeast-2 Region. All VPCs are attached to a transit gateway in ap-southeast-2. There are dedicated VPCs for each application service. The company also has VPCs for centralized security features such as proxies, firewalls, and logging.

The company plans to duplicate the infrastructure from ap-southeast-2 to the us-west-1 Region. A network engineer must establish connectivity between the various applications in the two Regions. The solution must maximize bandwidth, minimize latency and minimize operational overhead.

Which solution will meet these requirements?

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

- E. Create a VPN server in a VPC in each Region
- 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 130

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AWS-Certified-Advanced-Networking-Specialty Practice Exam Features:

- * AWS-Certified-Advanced-Networking-Specialty Questions and Answers Updated Frequently
- * AWS-Certified-Advanced-Networking-Specialty Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-Advanced-Networking-Specialty Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Certified-Advanced-Networking-Specialty Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Advanced-Networking-Specialty Practice Test Here](#)