PassQueen
a.k.a dumpsqueen.com

Amazon AWS Certified Advanced Networking - Specialty

Amazon AWS ANS-C01

Version Demo

Total Demo Questions: 43

Total Premium Questions: 438

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Design and implement hybrid IT network architectures at scale	165
Topic 2, Design and implement AWS networks	127
Topic 3, Design and implement for security and compliance	64
Topic 4, Manage, optimize, and troubleshoot the network	82
Total	438

QUESTION NO: 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. Publish the logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query and analyze the logs.

B. Enable NAT gateway access logs. Publish the logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query and analyze the logs.

C. Configure Traffic Mirroring on the NAT gateway's elastic network interface. Send the traffic to an additional EC2 instance. Use tools such as tcpdump and Wireshark to query and analyze the mirrored traffic.

D. Enable VPC flow logs on the NAT gateway's elastic network interface. Publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure. Use Athena to query and analyze the logs.

E. Enable NAT gateway access logs. Publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure. Use Athena to query and analyze the logs.

ANSWER: A D

Explanation:

Enabling VPC Flow Logs on the NAT gateway's elastic network interface and delivering the records to Amazon CloudWatch Logs provides a practical way to identify the private-source IP addresses, destination addresses, ports, protocols, byte counts, and accepted or rejected traffic associated with NAT gateway use. CloudWatch Logs Insights can then aggregate the records, for example by source address or destination address, to identify which workloads and external services account for the increase in processed bytes.

Delivering VPC Flow Logs for the NAT gateway elastic network interface to an Amazon S3 bucket is also appropriate. Flow log records in S3 can be queried with Amazon Athena after defining the log schema, enabling SQL-based investigation across a long retention period. Queries can rank source instances by transferred bytes, identify high-volume destinations, and examine the protocols or ports involved. VPC Flow Logs can be created for a network interface, and AWS specifically documents using flow logs to monitor traffic to and from a NAT gateway. These two destinations provide complementary operational approaches: CloudWatch Logs Insights for interactive log analysis and Athena for SQL queries over logs retained in S3.

See the Amazon VPC Flow Logs documentation and AWS guidance on NAT gateway monitoring and troubleshooting.

QUESTION NO: 2

A company has two VPCs in different AWS accounts that are connected by a VPC peering connection. The VPCs have non-overlapping CIDR blocks. An Amazon EC2 instance in each VPC has a public DNS hostname and a private IP address.

The company needs instances in one VPC to resolve the public DNS hostname of an instance in the peer VPC to that instance's private IP address. The company does not want to create and manage private hosted zones for this purpose.

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

A. Enable DNS resolution support for the requester VPC to resolve public DNS hostnames in the accepter VPC to private IP addresses.

B. Enable DNS resolution support for the accepter VPC to resolve public DNS hostnames in the requester VPC to private IP addresses.

C. Create an Amazon Route 53 Resolver inbound endpoint in each VPC and configure the endpoint IP addresses as DHCP name servers.

D. Create public hosted zones for both VPC CIDR blocks and add A records for all EC2 private IP addresses.

E. Enable transit gateway DNS support on the VPC peering connection.

ANSWER: A B

Explanation:

Enable DNS resolution support for the VPC peering connection from the requester VPC to the accepter VPC and from the accepter VPC to the requester VPC. When this option is enabled, instances in either VPC can resolve the public DNS hostname of an instance in the peer VPC to the instance's private IP address.

The VPCs must have DNS support and DNS hostnames enabled, and the relevant VPC route tables and security groups must allow the actual traffic over the peering connection. DNS resolution support is a peering-connection option that is configured separately for each direction. See the VPC peering DNS resolution documentation.

QUESTION NO: 3

A company has three VPCs in a single AWS Region. Each VPC contains 15 Amazon EC2 instances, and no connectivity exists between the VPCs.

The company is deploying a new application across all three VPCs. The application requires high bandwidth between the nodes. A network engineer must implement connectivity between the VPCs.

Which solution will meet these requirements with the HIGHEST throughput?

A. Configure a transit gateway. Attach each VPC to the transit gateway. Configure static routing in each VPC to route traffic to the transit gateway.

B. Configure VPC peering between the three VPCs. Configure static routing to route traffic between the three VPCs.

C. Configure a transit VPC.

D. Configure AWS Site-to-Site VPN connections between each VPC.

ANSWER: B

Explanation:

Configure VPC peering between the three VPCs. Configure static routing to route traffic between the three VPCs. is the correct solution because it creates direct, private network paths between every pair of VPCs in the same Region. For three VPCs, the engineer should establish a full mesh of peering connections and add routes in the relevant route tables for each remote VPC CIDR range. This enables every application node to communicate directly with nodes in either of the other VPCs.

Regional VPC peering uses the AWS global network rather than the public internet and does not impose an aggregate bandwidth limit on the peering connection itself. The achievable throughput is therefore governed primarily by the networking capacity of the participating EC2 instances and their placement, allowing the application to use the highest available bandwidth between nodes. AWS also specifies that VPC peering is nontransitive, which is why direct peering relationships between all three VPCs are necessary rather than relying on one VPC to forward traffic for another. See the Amazon VPC Peering Guide and AWS guidance for full-mesh VPC peering.

QUESTION NO: 4

A network engineer is working on a private DNS design to integrate AWS workloads and on-premises resources. The AWS deployment consists of five VPCs in the eu-west-1 Region that connect to the on-premises network over AWS Direct Connect. The VPCs communicate with each other by using a transit gateway. Each VPC is associated with a private hosted zone that uses the aws.example.internal domain. The network engineer creates an Amazon Route 53 Resolver outbound endpoint in a shared services VPC and attaches the shared services VPC to the transit gateway.

The network engineer is implementing a solution for DNS resolution. Queries for hostnames that end with aws.example.internal must use the private hosted zone. Queries for hostnames that end with all other domains must be forwarded to a private on-premises DNS resolver.

Which solution will meet these requirements?

A. Add a forwarding rule for""""that targets the on-premises server's DNS IP address. Add a system rule for aws.example.internal that targets Route 53 Resolver.

B. Add a forwarding rule for aws example.internal that targets Route 53 Resolver. Add a system rule for V that targets the Route 53 Resolver outbound endpoint.

C. Add a forwarding rule for""""that targets the Route 53 Resolver outbound endpoint.

D. Add a forwarding rule for"." that targets the Route 53 Resolver outbound endpoint.

ANSWER: D

Explanation:

Add a forwarding rule for "." that targets the Route 53 Resolver outbound endpoint. A Route 53 Resolver forwarding rule for the root domain, represented by a period, acts as a catch-all rule for DNS namespaces that do not have a more specific local resolution path. The rule should use the centralized outbound endpoint and specify the private IP addresses of the on-premises DNS resolver as its forwarding targets. The rule must also be associated with each VPC that requires this behavior; it can be shared across accounts by using AWS Resource Access Manager when appropriate.

For queries under `aws.example.internal`, the associated private hosted zone provides the more specific matching DNS namespace, so Resolver answers from that private hosted zone. For names outside that namespace, the root forwarding rule sends the query through the outbound endpoint to the on-premises resolver. The outbound endpoint is deployed in the shared-services VPC, and network routing and security controls must permit the endpoint ENIs to reach the on-premises resolver over Direct Connect. This design centralizes outbound DNS forwarding while retaining authoritative private hosted-zone resolution for AWS names. See the Amazon Route 53 Resolver outbound forwarding documentation and private hosted zone considerations.

QUESTION NO: 5

Company A recently acquired Company

B.

Company A has a hybrid AWS and on-premises environment that uses a hosted AWS Direct Connect connection, a Direct Connect gateway, and a transit gateway. Company A has a transit VIF to access the resources in itsproduction environment in the us-east-1 Region.

Company B has applications that run across multiple VPCs in the us-west-2 Region in a single AWS account. A transit gateway connects all Company B's application VPCs. The CIDR blocks for both companies do not overlap.

Company A needs to use the existing Direct Connect connection to access Company B’s applications from the on-premises environment.

Which solution will meet these requirements?

A. Create a new Direct Connect gateway in the Company B account. Associate the Company B transit gateway with the new Direct Connect gateway. Create a transit VIF on the existing hosted connection for Company B.

B. Create an association proposal from the Company B account to associate the Company B transit gateway with the Company A Direct Connect gateway. Accept the transit gateway association proposal by logging into the Company A account.

C. Create multiple virtual private gateways. Attach the virtual private gateways to each of Company B's application VPCs. Create a hosted private VIF for each virtual private gateway.

D. Create a new Direct Connect gateway in the Company B account. Associate the Company B transit gateway with the new Direct Connect gateway. Create a hosted private VIF for Company B.

ANSWER: B

Explanation:

Create an association proposal from the Company B account to associate the Company B transit gateway with the Company A Direct Connect gateway. Accept the transit gateway association proposal by logging into the Company A account. This uses the existing transit virtual interface and Direct Connect gateway, avoiding the need for another Direct Connect connection or additional virtual interfaces. A Direct Connect gateway can be associated with transit gateways, including a transit gateway owned by a different AWS account. For a cross-account association, the transit gateway owner creates an association proposal that identifies the Direct Connect gateway and permitted prefixes; the Direct Connect gateway owner then accepts that proposal.

After the association is accepted, the Direct Connect gateway provides the routing relationship between Company A's on-premises network, through its existing transit VIF, and Company B's transit gateway in us-west-2. Company B's transit gateway can then route traffic to its attached application VPCs. Because the companies' CIDR ranges do not overlap, the required prefixes can be advertised and allowed through the Direct Connect gateway association without address-space ambiguity. This architecture also supports the intended multi-Region connectivity model: Direct Connect gateway associations are not limited to the Region of the transit VIF or transit gateway. AWS documents the cross-account proposal workflow and routing considerations at Direct Connect gateways and Direct Connect gateway associations.

QUESTION NO: 6

A company hosts its IT infrastructure in an on-premises data center. The company wants to migrate the infrastructure to the AWS Cloud in phases. A network engineer wants to set up a 10 Gbps AWS Direct Connect dedicated connection between the on-premises data center and VPCs. The company's network provider needs 3 months to provision the Direct Connect connection.

In the meantime, the network engineer implements a temporary solution by deploying an AWS Site-to-Site VPN connection that terminates to a virtual private gateway. The network engineer observes that the bandwidth of the Site-to-Site VPN connection is capped at 1.25 Gbps despite a powerful customer gateway device.

What should the network engineer do to improve the VPN connection bandwidth before the implementation of the Direct Connect connection?

A. Contact AWS Support to request a bandwidth quota increase for the existing Site-to-Site VPN connection.

B. Discuss the issue with the hardware vendor. Buy a bigger and more powerful customer gateway device that has faster encryption and decryption capabilities.

C. Create several additional Site-to-Site VPN connections that terminate on the same virtual gateway. Configure equal-cost multi-path (ECMP) routing to use all the VPN connections simultaneously.

D. Create a transit gateway. Attach the VPCs to the transit gateway. Create several additional Site-to-Site VPN connections that terminate on the transit gateway. Configure equal-cost multi-path (ECMP) routing to use all the VPN connections simultaneously.

ANSWER: C

Explanation:

Create several additional Site-to-Site VPN connections that terminate on the same virtual gateway and configure equal-cost multi-path (ECMP) routing to use all the VPN connections simultaneously. AWS Site-to-Site VPN tunnel throughput is limited to up to 1.25 Gbps under ideal conditions, so a single VPN connection cannot provide the required aggregate throughput merely because the customer gateway has more processing capacity. AWS supports scaling VPN bandwidth horizontally by establishing multiple VPN connections to the same virtual private gateway and advertising identical routes with equal BGP attributes. ECMP then distributes eligible traffic flows across the equal-cost VPN paths, increasing aggregate available bandwidth while retaining the virtual private gateway architecture already deployed for the temporary migration solution. Each connection includes redundant tunnels, and the customer gateway must support BGP and ECMP so that route advertisements and traffic distribution are configured consistently. This approach provides a practical interim capacity increase until the dedicated Direct Connect connection is available. AWS documents the per-tunnel throughput expectation and recommends multiple VPN connections with ECMP when additional aggregate bandwidth is needed. See the AWS Site-to-Site VPN quotas and limits and AWS guidance for redundant VPN connections to a virtual private gateway.

QUESTION NO: 7

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. Connect the inspection VPC to the transit gateway by using a VPC attachment. Create a target group, and register the appliances with the target group. Create a Network Load Balancer (NLB), and set it up to forward to the newly created target group. Configure a default route in the inspection VPCs transit gateway subnet toward the NLB.

B. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VPC. Connect the inspection VPC to the transit gateway by using a VPC attachment. Create a target group, and register the appliances with the target group. Create a Gateway Load Balancer, and set it up to forward to the newly created target group. Configure a default route in the inspection VPC's transit gateway subnet toward the Gateway Load Balancer endpoint.

C. Configure two route tables on the transit gateway. Associate one route table with all the attachments of the application VPCs. Associate the other route table with the inspection VPC's attachment. Propagate all VPC attachments into the inspection route table. Define a static default route in the application route table. Enable appliance mode on the attachment that connects the inspection VPC.

D. Configure two route tables on the transit gateway. Associate one route table with all the attachments of the application VPCs. Associate the other route table with the inspection VPCs attachment. 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.

E. Configure one route table on the transit gateway. Associate the route table with all the VPCs. Propagate all VPC attachments into the route table. Define a static default route in the route table.

ANSWER: B C

Explanation:

Deploying the inspection appliances behind a Gateway Load Balancer provides a scalable and highly available service-insertion layer for third-party virtual appliances. Gateway Load Balancer distributes flows to healthy registered appliances, performs health checks, and maintains flow affinity so that packets for a flow are consistently sent to the same appliance. Gateway Load Balancer endpoints provide the transparent routing target in the inspection VPC, allowing traffic arriving from the transit gateway attachment subnet to be sent through the inspection fleet.

Using separate Transit Gateway route tables creates the required service-insertion path. The application-VPC attachment route table sends traffic to the inspection VPC through a static default route, while the inspection-VPC attachment route table learns the application VPC CIDRs through propagation and can return inspected traffic to its intended destination. Enabling appliance mode on the inspection VPC attachment is essential for stateful appliances that cannot support asymmetric routing. Appliance mode preserves Availability Zone affinity for flows through that attachment, ensuring that forward and return packets traverse the same inspection path while still supporting high availability across Availability Zones. AWS documents this centralized inspection pattern with Gateway Load Balancer and Transit Gateway appliance mode in its centralized inspection guidance and Transit Gateway appliance deployment documentation.

QUESTION NO: 8

A company has hundreds of VPCs on AWS. All the VPCs access the public endpoints of Amazon S3 and AWS Systems Manager through NAT gateways. All the traffic from the VPCs to Amazon S3 and Systems Manager travels through the NAT gateways. The company's network engineer must centralize access to these services and must eliminate the need to use public endpoints.

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

A. Create a central egress VPC that has private NAT gateways. Connect all the VPCs to the central egress VPC by using AWS Transit Gateway. Use the private NAT gateways to connect to Amazon S3 and Systems Manager by using private IP addresses.

B. Create a central shared services VPC. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access. Ensure that private DNS is turned off. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway. Create an Amazon Route 53 forwarding rule for each interface VPC endpoint. Associate the forwarding rules with all the VPCs. Forward DNS queries to the interface VPC endpoints.

C. Create a central shared services VPC. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager. Ensure that private DNS is turned off. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway. Create Amazon Route 53 private hosted zones for the full service endpoint names, associate the private hosted zones with all the VPCs, and create alias records that point to the interface VPC endpoints.

D. Create a central shared services VPC. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway. Ensure that private DNS is turned on for the interface VPC endpoints and that the transit gateway is created with DNS support turned on.

ANSWER: C

Explanation:

Create a central shared services VPC, deploy interface VPC endpoints for Amazon S3 and AWS Systems Manager there, disable private DNS on those endpoints, and use Amazon Route 53 private hosted zones associated with the connected VPCs. The private hosted zones should use each AWS service's regional public DNS namespace and contain alias records that direct the service name to the corresponding interface endpoint. This causes workloads in every associated VPC to resolve normal AWS service names to the private endpoint ENIs in the shared services VPC. Traffic can then traverse AWS Transit Gateway privately to reach those endpoint ENIs, without using NAT gateways or public service endpoints.

This design centralizes endpoint provisioning, endpoint policy administration, and DNS management rather than requiring endpoints in hundreds of individual VPCs. Disabling private DNS is important because private DNS for an interface endpoint applies only within the VPC that owns the endpoint; the shared private hosted zone provides the required cross-VPC DNS resolution. Route 53 private hosted zones can be associated with multiple VPCs, including VPCs in different AWS accounts through authorization and association workflows. AWS documents centralized interface endpoint access using private hosted zones and alias records, and lists both Amazon S3 and Systems Manager as services supporting interface endpoints. See Centralized access to VPC private endpoints and AWS services that integrate with AWS PrivateLink.

QUESTION NO: 9

A company has an AWS environment that includes multiple VPCs that are connected by a transit gateway. The company has decided to use AWS Site-to-Site VPN to establish connectivity between its on-premises network and its AWS environment.

The company does not have a static public IP address for its on-premises network. A network engineer must implement a solution to initiate the VPN connection on the AWSside of the connection for traffic from the AWS environment to the on-premises network.

Which combination of steps should the network engineer take to establish VPN connectivity between the transit gateway and the on-premises network? (Choose three.)

A. Configure the Site-to-Site VPN tunnel options to use Internet Key Exchange version 1 (IKEv1).

B. Configure the Site-to-Site VPN tunnel options to use Internet Key Exchange version 2 (IKEv2).

C. Use a private certificate authority (CA) from AWS Private Certificate Authority to create a certificate.

D. Use a public certificate authority (CA) from AWS Private Certificate Authority to create a certificate.

E. Create a customer gateway. Specify the current dynamic IP address of the customer gateway device's external interface.

F. Create a customer gateway and specify 0.0.0.0 as the IP address of the customer gateway device.

ANSWER: B C F

Explanation:

For a customer gateway device whose public-facing address changes dynamically, AWS Site-to-Site VPN supports certificate-based authentication rather than relying on a fixed peer IP address. This configuration requires Internet Key Exchange version 2 (IKEv2), which supports the certificate-based authentication workflow for a dynamically addressed customer gateway. The customer gateway resource must use `0.0.0.0` as its IP address; this explicitly identifies the peer as dynamically addressed and allows the AWS VPN endpoint to establish the tunnel without a preconfigured fixed on-premises address.

The customer gateway device needs a certificate to authenticate during IKEv2 negotiation. A certificate issued by a private certificate authority in AWS Private Certificate Authority is appropriate because the organization can establish and control the trust chain used by its AWS VPN endpoint and on-premises device. The certificate and its issuing certificate authority are then used in the certificate-based VPN configuration. Together, IKEv2, a private-CA-issued certificate, and a customer gateway configured with `0.0.0.0` provide the AWS-supported design for this dynamic-IP requirement. See the AWS Site-to-Site VPN customer gateway options documentation and the AWS documentation for certificate-based Site-to-Site VPN authentication.

QUESTION NO: 10

A company has deployed IPv6-only Amazon EC2 instances in private subnets. The instances must connect to an external software repository that provides only IPv4 endpoints. The company does not want to assign public IPv4 addresses to the instances or deploy a proxy fleet.

Which solution will provide the required connectivity?

A. Enable DNS64 for the private subnets. Deploy a NAT gateway, and add a route for 64:ff9b::/96 that points to the NAT gateway.

B. Create an egress-only internet gateway, and add a route for 0.0.0.0/0 that points to the egress-only internet gateway.

C. Deploy a private NAT gateway, and add a route for ::/0 that points to the private NAT gateway.

D. Associate an IPv4 CIDR block with the private subnets, and configure an internet gateway as the default route target.

ANSWER: A

Explanation:

Enable DNS64 for the private subnets, deploy a NAT gateway, and configure a route for the well-known NAT64 prefix, `64:ff9b::/96`, to the NAT gateway. DNS64 synthesizes IPv6 addresses from IPv4 DNS responses when an IPv6-only instance queries a hostname that has no AAAA record. The instance then sends traffic to the synthesized IPv6 destination.

The NAT gateway performs NAT64 translation for traffic addressed to the NAT64 prefix. It translates the IPv6 packets to IPv4 packets and forwards them to the IPv4 destination. This allows IPv6-only workloads to initiate connections to IPv4 services without requiring public IPv4 addresses on the workloads. See NAT gateways, NAT64, and DNS64.

QUESTION NO: 11

A company has a 10 Gbps AWS Direct Connect dedicated connection to a transit gateway. The company must encrypt traffic between its on-premises data center and AWS without reducing the available bandwidth. The company also requires a backup encrypted connection that can use the existing Direct Connect connection if the primary Direct Connect connection fails.

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

A. Configure MACsec on the dedicated Direct Connect connection and on compatible customer edge equipment.

B. Create a public VIF on the Direct Connect connection.

C. Create an AWS Site-to-Site VPN connection to the transit gateway and configure dynamic routing.

D. Create a private VIF and use the private VIF to establish an AWS Site-to-Site VPN connection to the transit gateway.

E. Create a NAT gateway in the transit gateway attachment subnet to encrypt traffic that traverses the Direct Connect connection.

F. Configure an egress-only internet gateway as the backup path for the transit gateway attachment.

ANSWER: A B C

Explanation:

MACsec provides Layer 2 encryption on supported dedicated Direct Connect connections and preserves the high throughput of the physical connection. Configure MACsec on the dedicated Direct Connect connection and on compatible customer edge equipment to encrypt the primary path.

For the backup path, create an AWS Site-to-Site VPN connection that terminates on the transit gateway. A Site-to-Site VPN connection provides IPsec encryption and can use the public IP addresses of AWS VPN endpoints. Create a public VIF on the Direct Connect connection so that the customer gateway can reach those public VPN endpoints through Direct Connect instead of through the public internet. Dynamic BGP routing allows the VPN routes to be used as a backup when the Direct Connect transit VIF path is unavailable.

See AWS Direct Connect MACsec encryption, AWS Direct Connect virtual interfaces, and AWS Site-to-Site VPN.

QUESTION NO: 12

A retail company is migrating its on-premises application to the AWS Cloud. Currently, the company has two on-premises data center locations. One data center is on the east coast of the United States, and one data center is on the west coast.

Each data center hosts four database systems. The largest database system stores 500 GB of data. The data centers are interconnected by two 10 GbE circuits for data synchronization. Each data center has two separate 1 GbE upstream internet connections. The company plans to have eight total VPCs to service its multiple business units. Four VPCs will be in the us-east-1 Region, and four will be in the us-west-2 Region.

A network engineer needs to design a connectivity solution that allows VPC-to-VPC connectivity. The solution must also allow secure connections between the on-premises data centers and AWS during the migration process. The company expects spikes in traffic among the VPCs during database synchronization. The company wants to run themigration plan during one weekend and as soon as technically possible. The company also wants to minimize long-term operational and human resources costs.

Which combination of steps will meet these requirements? {Select TWO.)

A. Deploy a transit gateway in each Region, attach the local VPCs to it, and peer the transit gateways. Update the transit gateway and VPC route tables to allow any VPC to connect to any other VPC.

B. Configure VPC peering between all the VPCs. Update the VPC route tables to allow connectivity.

C. Provision two AWS Direct Connect connections from two Direct Connect locations that serve us-east-1 and us-west-2 to provide connectivity between the data centers and AWS.

D. Provision one transit gateway VPN attachment for each data center to build connectivity between the on-premises data centers and AWS VPCs.

E. Provision one AWS Site-to-Site VPN connection for each data center and for each VPC to build connectivity between the on-premises data centers and AWS VPCs.

ANSWER: A C

Explanation:

Deploying a transit gateway in each Region, attaching the local VPCs, and peering the transit gateways provides scalable many-to-many connectivity without maintaining a full mesh of VPC peering relationships. Transit Gateway is a regional

service, so the VPCs in each Region attach to their local transit gateway. A transit gateway peering attachment then carries traffic between the Regional transit gateways. Route tables on the transit gateways and VPCs control which CIDR ranges are reachable. This design is operationally simpler as the number of VPCs grows and is well suited to bursty inter-VPC database synchronization traffic.

Two AWS Direct Connect connections, using Direct Connect locations serving the two Regions, provide dedicated private network connectivity from the data centers to AWS. Direct Connect can deliver substantially more predictable throughput than the existing internet links, supporting the requirement to transfer migration data quickly. A Direct Connect gateway can associate the private virtual interfaces with the transit gateways so the on-premises networks can reach the VPCs through the centralized routing design. For encrypted traffic, the company can additionally use a Site-to-Site VPN over Direct Connect. See the AWS Transit Gateway peering documentation and the AWS Direct Connect gateway documentation.

QUESTION NO: 13

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

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

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

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

B. Create a certificate for service.example.com by using AWS Certificate Manager (ACM). Configure CloudFront to use this custom SSL/TLS certificate. Change the default behavior to redirect HTTP to HTTPS.

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

D. Create a public certificate from a third-party certificate provider with any domain name for the EC2 instances. Configure the backend to use this certificate for its HTTPS listener. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its targets. Attach the existing Auto Scaling group to this new target group.

E. Create a certificate for service-alb.example.com by using AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the service-alb.example.com ACM certificate. Modify the CloudFront origin to use the HTTPS protocol only. Delete the HTTP listener on the ALB.

F. Create a self-signed certificate for service-alb.example.com. Import the certificate into AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the imported service-alb.example.com ACM certificate. Modify the CloudFront origin to use the HTTPS protocol only. Delete the HTTP listener on the ALB.

ANSWER: B D E

Explanation:

End-to-end encryption requires TLS on each network segment: from viewers to CloudFront, from CloudFront to the ALB origin, and from the ALB to the Windows EC2 targets. Creating an ACM certificate for `service.example.com`, associating it with the CloudFront distribution, and redirecting HTTP requests to HTTPS encrypts viewer traffic to the CloudFront edge. For CloudFront, the ACM certificate used with an alternate domain name must cover that domain and be requested or imported in the US East (N. Virginia) Region.

Creating an ACM certificate for `service-alb.example.com`, configuring an ALB HTTPS listener with that certificate, and setting the CloudFront origin protocol policy to HTTPS Only ensures CloudFront uses TLS when retrieving content from the origin. Removing the ALB HTTP listener prevents unencrypted origin requests.

Finally, configuring the Windows instances to serve HTTPS with a third-party certificate and using an HTTPS target group ensures that the ALB establishes TLS connections to registered EC2 targets. ALB target groups support HTTPS health checks and HTTPS traffic to instance targets; the load balancer encrypts traffic to those targets. Together, these changes

preserve encryption continuously from the client through every hop to the application backend. See CloudFront HTTPS configuration and Application Load Balancer target groups.

QUESTION NO: 14

A company runs a TCP application on Amazon EC2 instances in three Availability Zones. The instances are registered in a Network Load Balancer (NLB) target group. The NLB has subnet mappings in all three Availability Zones, but an application deployment temporarily removes all healthy targets from one Availability Zone.

Clients use the NLB static IP address in each Availability Zone. The company needs clients that connect to the static IP address in the Availability Zone without healthy targets to be served by healthy targets in the other Availability Zones.

Which solution will meet these requirements?

A. Enable cross-zone load balancing on the Network Load Balancer.

B. Change the Network Load Balancer listener from TCP to TLS.

C. Configure an Application Load Balancer as a target of the Network Load Balancer.

D. Create an interface VPC endpoint in each Availability Zone and configure the clients to use the endpoint addresses.

ANSWER: A

Explanation:

Enable cross-zone load balancing on the Network Load Balancer. With cross-zone load balancing enabled, an NLB node can distribute traffic to healthy targets in all enabled Availability Zones, rather than distributing traffic only to targets in its local Availability Zone. Therefore, a client that connects through an NLB static IP address in an Availability Zone that has no healthy local targets can still reach a healthy target in another Availability Zone.

Cross-zone load balancing is configured at the load balancer level for an NLB. The target group health checks must continue to identify healthy application instances, and normal security group and network ACL requirements still apply. See the Network Load Balancer cross-zone load balancing documentation.

QUESTION NO: 15

A company has its production VPC (VPC-A) in the eu-west-1 Region in Account 1. VPC-A is attached to a transit gateway (TGW-A) that is connected to an on-premises data center in Dublin, Ireland, by an AWS Direct Connect transit VIF that is configured for an AWS Direct Connect gateway. The company also has a staging VPC (VPC-B) that is attached to another transit gateway (TGW-B) in the eu-west-2 Region in Account 2.

A network engineer must implement connectivity between VPC-B and the on-premises data center in Dublin.

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

A. Configure inter-Region VPC peering between VPC-A and VPC-B. Add the required VPC peering routes. Add the VPC-B CIDR block in the allowed prefixes on the Direct Connect gateway association.

B. Associate TGW-B with the Direct Connect gateway. Advertise the VPC-B CIDR block under the allowed prefixes.

C. Configure another transit VIF on the Direct Connect connection and associate TGW-B. Advertise the VPC-B CIDR block under the allowed prefixes.

D. Configure inter-Region transit gateway peering between TGW-A and TGW-B. Add the peering routes in the transit gateway route tables. Add both the VPC-A and the VPC-B CIDR block under the allowed prefix list in the Direct Connect gateway association.

E. Configure an AWS Site-to-Site VPN connection over the transit VIF to TGW-B as a VPN attachment.

ANSWER: B C

Explanation:

Associate TGW-B with the Direct Connect gateway. Advertise the VPC-B CIDR block under the allowed prefixes. This is a supported design because a Direct Connect gateway can be associated with multiple transit gateways, including transit gateways in different AWS Regions and AWS accounts. The Direct Connect gateway provides the routing relationship between the transit VIF and each associated transit gateway. The allowed-prefix list for the TGW-B association must include VPC-B's CIDR range so that the on-premises network can receive and route traffic destined for that VPC. Appropriate transit gateway and VPC route-table entries are also required for end-to-end forwarding.

Configure another transit VIF on the Direct Connect connection and associate TGW-B. Advertise the VPC-B CIDR block under the allowed prefixes. A separate transit VIF can likewise be used to provide Direct Connect connectivity for TGW-B through a Direct Connect gateway association. This provides an additional valid implementation path when the organization wants to use a distinct transit VIF for the staging connectivity. As with any transit gateway association to a Direct Connect gateway, VPC-B's prefix must be authorized in the allowed-prefix list and advertised to the on-premises side. AWS documents Direct Connect gateway and transit gateway associations at AWS Direct Connect gateways and the associated prefix requirements at Allowed prefixes for Direct Connect gateway associations.

QUESTION NO: 16

A development team is building a new web application in the AWS Cloud. The main company domain, example.com. is currently hosted in an Amazon Route 53 public hosted zone in one of the company's production AWS accounts.

The developers want to test the web application in the company's staging AWS account by using publicly resolvable subdomains under the example.com domain with the ability to create and delete DNS records as needed. Developers have full access to Route 53 hosted zones within the staging account, but they are prohibited from accessing resources in any of the production AWS accounts.

Which combination of steps should a network engineer take to allow the developers to create records under the example.com domain? (Select TWO.)

A. Create a public hosted zone for example.com in the staging account.

B. Create a staging.example.com NS record in the example.com domain. Populate the value with the name servers from the staging.example.com domain. Set the routing policy type to simple routing.

C. Create a private hosted zone for stagmg.example.com in the staging account.

D. Create an example.com NS record in the staging.example.com domain. Populate the value with the name servers from the example.com domain. Set the routing policy type to simple routing

E. Create a public hosted zone for staging.example.com in the staging account.

ANSWER: B E

Explanation:

The required design is DNS subdomain delegation. First, create a public hosted zone for `staging.example.com` in the staging account. Route 53 assigns that hosted zone a set of authoritative name servers and automatically creates its apex NS and SOA records. Because the hosted zone is in the staging account, developers can use their existing permissions there to independently add, update, and delete records such as `app.staging.example.com`, while those records remain publicly resolvable.

Next, a network engineer with access to the production account must create an NS record named `staging.example.com` in the parent `example.com` public hosted zone. The NS record values must be the four Route 53 name servers assigned to the staging public hosted zone. This delegates authority for the entire `staging.example.com` namespace from the production zone to the staging zone. Public recursive resolvers can then follow the delegation from `example.com` to the staging account's authoritative name servers. A simple routing policy is appropriate for this delegation record. AWS documents the process of delegating a subdomain in Routing traffic for subdomains and explains public hosted zones in the Route 53 hosted zone documentation.

QUESTION NO: 17

A company is planning to host a secure web application across multiple Amazon EC2 instances. The application will have an associated DNS domain in an Amazon Route 53 hosted zone.

The company wants to protect the domain from DNS poisoning attacks. The company also wants to allow web browsers to authenticate into the application by using a trusted third party.

Which combination of actions will meet these requirements?

A. Configure the Route 53 hosted zone to use DNS Security Extensions (DNSSEC). Install self-signed X.509 certificates on the EC2 instances.

B. Configure a Name Authority Pointer (NAPTR) record in the Route 53 hosted zone. Install X 509 certificates that are signed by a public certificate authority on the EC2 instances.

C. Configure the Route 53 hosted zone to use DNS Security Extensions (DNSSEC). Install X.509 certificates that are signed by a public certificate authority on the EC2 instances.

D. Configure a Name Authority Pointer (NAPTR) record in the Route 53 hosted zone. Install self-signed X.509 certificates on the EC2 instances.

ANSWER: C

Explanation:

Configure the Route 53 hosted zone to use DNS Security Extensions (DNSSEC). Install X.509 certificates that are signed by a public certificate authority on the EC2 instances. DNSSEC protects the integrity and authenticity of DNS responses by digitally signing DNS records. A validating recursive DNS resolver can verify the chain of trust from the DNS root through the parent zone to the Route 53 hosted zone. This validation helps clients detect forged or modified DNS responses, addressing DNS cache-poisoning risks for the domain. Route 53 supports DNSSEC signing for public hosted zones, with a key-signing key managed through AWS Key Management Service.

For browser-based HTTPS authentication, the servers need certificates that browsers trust. An X.509 TLS certificate issued by a publicly trusted certificate authority provides that third-party trust relationship: the browser validates the certificate chain, the domain name in the certificate, and certificate validity before establishing the encrypted connection. Although the question describes certificates installed on EC2 instances, in a production multi-instance design the same public certificate can also be deployed through a load balancer or other TLS termination point. See the AWS documentation for configuring DNSSEC signing in Route 53 and the AWS Certificate Manager guidance on public certificates and browser trust.

QUESTION NO: 18

A company has its production VPC (VPC-A) in the eu-west-1 Region in Account 1. VPC-A is attached to a transit gateway (TGW-A) that is connected to an on-premises data center in Dublin, Ireland, by an AWS Direct Connect transit VIF that is configured for an AWS Direct Connect gateway. The company also has a staging VPC (VPC-B) that is attached to another transit gateway (TGW-B) in the eu-west-2 Region in Account 2.

A network engineer must implement connectivity between VPC-B and the on-premises data center in Dublin.

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

A. Configure inter-Region VPC peering between VPC-A and VPC-B. Add the required VPC peering routes. Add the VPC-B CIDR block in the allowed prefixes on the Direct Connect gateway association.

B. Associate TGW-B with the Direct Connect gateway. Advertise the VPC-B CIDR block under the allowed prefixes.

C. Configure another transit VIF on the Direct Connect connection. Associate TGW-B with the Direct Connect gateway, and advertise the VPC-B CIDR block under the allowed prefixes.

D. Configure inter-Region transit gateway peering between TGW-A and TGW-B. Add the peering routes in the transit gateway route tables. Add both the VPC-A and the VPC-B CIDR block under the allowed prefix list in the Direct Connect gateway association.

E. Configure an AWS Site-to-Site VPN connection over the transit VIF to TGW-B as a VPN attachment.

ANSWER: B C

Explanation:

“Associate TGW-B with the Direct Connect gateway. Advertise the VPC-B CIDR block under the allowed prefixes.” is valid because an AWS Direct Connect gateway can associate with multiple transit gateways, including transit gateways in different Regions and accounts. The association for TGW-B must include the VPC-B CIDR range as an allowed prefix. This allows AWS to advertise the reachable VPC-B prefix to the on-premises network over BGP and allows traffic destined for the data center to use the transit VIF and Direct Connect gateway.

“Configure another transit VIF on the Direct Connect connection. Associate TGW-B with the Direct Connect gateway, and advertise the VPC-B CIDR block under the allowed prefixes.” is also valid. A transit virtual interface connects the Direct Connect connection to a Direct Connect gateway, and the gateway association makes the transit gateway’s attached networks reachable. Using an additional transit VIF provides another logical Direct Connect connectivity path to the same Direct Connect gateway while the TGW-B association and allowed-prefix configuration establish the required routing for VPC-B. The on-premises router must accept and advertise the corresponding BGP routes so traffic can flow bidirectionally.

See the AWS documentation for Direct Connect gateways and transit virtual interfaces and transit gateway associations.

QUESTION NO: 19

A company has a VPC that includes application workloads that run on Amazon EC2 instances in a single AWS Region. The company wants to use AWS Local Zones to deploy an extension of the application workloads that run in the Region. The extended workloads in the Local Zone need to communicate bidirectionally with the workloads in the VPC in the Region.

Which solution will meet these requirements MOST cost-effectively?

A. Create a new VPC in the Local Zone. Attach all the VPCs to a transit gateway. Configure routing for the transit gateway and the VPCs. Deploy instances in the new VP

B. Deploy a third-party appliance in a new VPC in the Region. Create a new VPC in the Local Zone. Create VPN connections to the appliance for the VPCs. Deploy instances in the new VPC in the Local Zone.

C. Create a new subnet in the Local Zone. Deploy a third-party appliance in the VPC with interfaces in each subnet. Configure the new subnet to route the Local Zone through the appliance. Deploy instances in the new subnet.

D. Create a new subnet in the Local Zone. Configure the new subnet to use a CIDR block that is within the VPC’s CIDR block. Deploy instances in the new subnet in the Local Zone.

ANSWER: D

Explanation:

Create a new subnet in the Local Zone. Configure the new subnet to use a CIDR block that is within the VPC’s CIDR block. Deploy instances in the new subnet in the Local Zone. This is the most cost-effective approach because a Local Zone is an extension of its parent AWS Region, and Local Zone subnets are created directly in the existing VPC. The EC2 instances launched into that subnet therefore use the same VPC networking construct as the Regional workloads, including VPC route tables, security groups, and network ACLs. Traffic between the Local Zone subnet and subnets in the parent Region can be routed within the VPC, supporting the required bidirectional communication without introducing separate VPC connectivity infrastructure or network appliances. AWS requires a Local Zone subnet to receive an IPv4 CIDR range from the parent VPC’s CIDR block, so assigning an appropriate unused range and launching the extension workloads there directly meets the addressing requirement. This design avoids additional recurring charges and operational overhead associated with transit gateways, VPN connections, or third-party virtual appliances. For implementation details, see the AWS Local Zones for Amazon VPC documentation and the Amazon EC2 Local Zones guide.

QUESTION NO: 20

A media company is planning to host an event that the company will live stream to users. The company wants to use Amazon CloudFront.

A network engineer creates a primary origin and a secondary origin for CloudFront. Theengineer needs to ensure that the primary origin can fail over to the secondary origin within 15 seconds if a disruption occurs.

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

A. Configure a Lambda@Edge function to check the health status of both origins every 10 seconds. Reroute incoming requests when the origin health status is unhealthy.

B. Create a Network Load Balancer (NLB) in front of both origins Configure the NLB as the origin in CloudFront.

C. Set the CloudFront origin connection timeout value to 5 seconds Set the origin connection attempts value to 2.

D. Configure a Lambda@Edge function to monitor incoming requests for an origin response. Reroute incoming requests if no response is received from the primary origin within 10 seconds.

ANSWER: C

Explanation:

Set the CloudFront origin connection timeout value to 5 seconds Set the origin connection attempts value to 2. is the lowest-overhead configuration because CloudFront natively manages origin failover when the origins are configured in an origin group. For connection failures, CloudFront waits for the configured connection timeout and retries the primary origin up to the configured number of connection attempts before failing over to the secondary origin. A 5-second timeout with two attempts limits this connection phase to 10 seconds, which is within the required 15-second failover window. No custom health-monitoring code, edge-function deployment, or separate load-balancing infrastructure is required. The primary and secondary origins should be configured as a CloudFront origin group, with the desired failover criteria defined for the group. CloudFront then directs requests to the primary origin under normal conditions and uses the secondary origin when the configured failover conditions occur. AWS documents that the origin connection timeout is configurable from 1 to 10 seconds and that CloudFront can make from 1 to 3 connection attempts. See CloudFront origin settings and CloudFront origin failover.

QUESTION NO: 21

A company is replatforming a legacy data processing solution to AWS. The company deploys the solution on Amazon EC2 Instances in private subnets that are in one VPC.

The solution uses Amazon S3 for abject storage. Both the data that the solution processes and the data the solution produces are stored in Amazon S3. The solution uses Amazon DynamoDB to save its own state. The company collects flow logs for the VPC. Thesolution uses one NAT gateway to register its license through the internet. A software vendor provides a specific hostname so the solution can register its license.

The company notices that the AWS bill exceeds the projected budget for the solution. A network engineer uses AWS Cost Explorer to investigate the bill. The network engineer notices that the USE2-NatGateway-Bytes($) usage type is the root cause of the higher than expected bill.

What should the network engineer do to resolve the issue? (Choose two.)

A. Set up Amazon VPC Traffic Mirroring. Analyze the traffic to identify the traffic that the NAT gateway processes.

B. Examine the VPC flow logs to identity the traffic that traverses the NAT gateway.

C. Set up an AWS Cost and Usage Report in the AWS Billing and Cost Management console. Examine the report to find more details about the NAT gateway charges.

D. Verify that the security groups attached to the EC2 instances allow outgoing traffic only to the IP addresses that the hostname resolves to, the VPC CIDR block, and the AWS IP address ranges for Amazon S3 and DynamoDB.

E. Verify that the gateway VPC endpoints for Amazon S3 and DynamoDB are both set up and associated with the route tables of the private subnets.

ANSWER: B E

Explanation:

Examine the VPC flow logs to identity the traffic that traverses the NAT gateway. Flow logs provide network-flow records containing source and destination addresses, ports, protocol, byte counts, and accept or reject status. Reviewing these records enables the engineer to identify which private-subnet workloads and destinations account for the high volume of NAT gateway data processing. This is particularly useful for confirming whether S3 or DynamoDB traffic is incorrectly taking the default NAT gateway route rather than remaining on the AWS network.

Verify that the gateway VPC endpoints for Amazon S3 and DynamoDB are both set up and associated with the route tables of the private subnets. Gateway endpoints provide private route-table targets for Amazon S3 and DynamoDB. When the endpoint routes are associated with the EC2 private-subnet route tables, traffic to these services bypasses the NAT gateway. This removes NAT gateway data-processing charges for the application's potentially large object-storage transfers and DynamoDB state operations, while preserving the NAT gateway only for the required internet-based license registration. AWS documents that gateway endpoints have no additional charge, whereas NAT gateways charge for data processed. See Amazon S3 gateway endpoints and VPC Flow Logs.

QUESTION NO: 22

A company is growing rapidly. Data transfers between the company's on-premises systems and Amazon EC2 instances that run in VPCs are limited by the throughput of a single AWS Site-to-Site VPN connection between the company's on-premises data center firewall and an AWS Transit Gateway.

A network engineer must resolve the throttling by designing a solution that is highly available and secure. The solution also must scale the VPN throughput from on premises to the VPC resources to support the increase in traffic.

Which solution will meet these requirements?

A. Configure multiple dynamic BGP-based Site-to-Site VPN connections to the transit gateway Configure equal-cost multi-path routing

B. Configure multiple static routing-based Site-to-Site VPN connections to the transit gateway. Configure equal-cost multi-path routing

C. Configure a new Site-to-Site VPN connection to the transit gateway Enable acceleration for the Site-to-Site VPN connection.

D. Configure a software appliance-based VPN connection over the internet from the on-premises firewall to an EC2 instance that has a large instance size and networking capabilities.

ANSWER: A

Explanation:

Configure multiple dynamic BGP-based Site-to-Site VPN connections to the transit gateway Configure equal-cost multi-path routing meets the availability, security, and throughput requirements. Each AWS Site-to-Site VPN connection provides two redundant IPsec tunnels, and dynamic routing with BGP allows the on-premises network and AWS Transit Gateway to exchange routes and detect tunnel-path failures automatically. Creating multiple VPN connections provides multiple encrypted paths between the customer gateway and the transit gateway, avoiding dependence on the throughput capacity of one VPN connection.

AWS Transit Gateway supports equal-cost multipath routing for Site-to-Site VPN attachments that use dynamic routing. With ECMP enabled, traffic flows can be distributed across VPN tunnels and across multiple eligible VPN connections that advertise equal-cost routes. This aggregates available VPN bandwidth for multiple flows while retaining resilient path selection if an individual tunnel or connection becomes unavailable. The on-premises firewall or VPN device must likewise support BGP and ECMP to use the parallel paths effectively. AWS documents ECMP support and its dynamic-routing requirement in the Transit Gateway ECMP documentation. AWS also describes VPN tunnel redundancy and BGP-based dynamic routing in the AWS Site-to-Site VPN documentation.

QUESTION NO: 23

A company has two VPCs that use overlapping CIDR blocks. Each VPC contains private Amazon EC2 instances that must initiate connections to servers in an on-premises data center through an AWS Transit Gateway and AWS Direct Connect.

The on-premises firewall requires a unique source CIDR range for each VPC. The on-premises network must not be able to initiate unsolicited connections to the EC2 instances.

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

A. Deploy a private NAT gateway in each VPC. Use a unique, non-overlapping private CIDR range for the translated source addresses in each VPC.

B. Route on-premises destination traffic from the application subnets to the private NAT gateways. Route traffic from the NAT gateway subnets to the transit gateway, and configure return routes for the translated CIDR ranges.

C. Create a VPC peering connection between the two VPCs. Add routes for the overlapping CIDR block through the peering connection.

D. Configure the on-premises firewall to allow the overlapping VPC CIDR block as a source for both VPCs.

E. Deploy internet-facing NAT gateways in the VPCs and route on-premises traffic through an internet gateway.

ANSWER: A B

Explanation:

Deploy a private NAT gateway in each VPC and assign each private NAT gateway an address from a unique, non-overlapping CIDR range. A private NAT gateway performs source network address translation for traffic to private networks, including networks reached through a transit gateway and Direct Connect. The on-premises firewall can therefore use the translated source addresses to distinguish traffic that originated in each overlapping VPC.

Configure the application subnet route tables to direct on-premises destination traffic to the private NAT gateway. Configure the NAT gateway subnet route tables to direct the on-premises destinations to the transit gateway, and configure the return routes for the translated CIDR ranges through the appropriate transit gateway attachments. A NAT gateway is stateful, so return traffic for connections initiated by the EC2 instances is permitted while unsolicited inbound connections are not. See the Amazon VPC NAT gateway documentation and the NAT gateway routing scenarios documentation.

QUESTION NO: 24

A global film production company uses the AWS Cloud to encode and store its videocontent before distribution. The company's three global offices are connected to the us-east-1 Region through AWS Site-to-Site VPN links that terminate on a transit gateway with BGP routing activated.

The company recently started to produce content at a higher resolution to support 8K streaming. The size of the content files has increased to three times the size of the content files from the previous format. Uploads of files to Amazon EC2 instances are taking 10 times longer than they did with the previous format.

Which actions should a network engineer recommend to reduce the upload times? (Choose two.)

A. Create a second VPN tunnel from each office location to the transit gateway. Activate equal-cost multi-path (ECMP) routing.

B. Modify the transit gateway to activate Jumbo MTU on the VPN tunnels to each office location.

C. Replace the existing VPN tunnels with new tunnels that have acceleration activated.

D. Upgrade each EC2 instance to a modern instance type. Activate Jumbo MTU in the operating system.

E. Replace the existing VPN tunnels with new tunnels that have IGMP activated.

ANSWER: A C

Explanation:

Creating a second VPN tunnel from each office location to the transit gateway and activating equal-cost multi-path (ECMP) routing increases the available aggregate bandwidth for the VPN connection. Transit Gateway supports ECMP for Site-to-Site VPN attachments when dynamic routing is used. Because BGP is already enabled, the network can advertise equal-

cost paths and distribute eligible traffic flows across multiple active tunnels. This is particularly appropriate for large parallel file-transfer workloads, where multiple flows can take advantage of the additional tunnel capacity.

Replacing the existing VPN tunnels with new tunnels that have acceleration activated can also reduce transfer time. AWS Site-to-Site VPN acceleration uses the AWS Global Accelerator network to route VPN traffic through nearby AWS edge locations and then across the AWS global network. This can improve performance for offices that are geographically distant from the VPN endpoint Region or that experience inconsistent public-internet paths. The combination of accelerated VPN connectivity and ECMP provides both improved path performance and greater aggregate throughput, directly addressing the sharp increase in upload duration caused by substantially larger video files.

For implementation details, see the AWS Site-to-Site VPN acceleration documentation and the Transit Gateway Site-to-Site VPN documentation.

QUESTION NO: 25

A company is planning to use an AWS Transit Gateway hub and spoke architecture to migrate to AWS. The current on-premises multi-protocol label switching (MPLS) network has strict controls that enforce network segmentation by using MPLS VPNs. The company has provisioned two 10 Gbps AWS Direct Connect connections to provide resilient, high-speed, low-latency connectivity to AWS.

A security engineer needs to apply the concept of network segmentation to the AWS environment to ensure that virtual routing and forwarding (VRF) is logically separated for each of the company's software development environments. The number of MPLS VPNs will increase in the future. On-premises MPLS VPNs will have overlapping address space. The company's AWS network design must support overlapping address space for the VPNs.

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

A. Deploy a software-defined WAN (SD-WAN) head-end virtual appliance and an SD-WAN controller into a Transit Gateway Connect VPC. Configure the company's edge routers to be managed by the new SD-WAN controller and to use SD-WAN to segment the traffic into the defined segments for each of the company's development environments.

B. Configure IPsec VPNs on the company edge routers for each MPLS VPN for each of the company's development environments. Attach each IPsec VPN tunnel to a discrete MPLS VPN. Configure AWS Site-to-Site VPN connections that terminate at a transit gateway for each MPLS VPN. Configure a transit gateway route table that matches the MPLS VPN for each Transit Gateway VPN attachment.

C. Create a transit VPC that terminates at the AWS Site-to-Site VRF-aware IPsec VPN. Configure IPsec VPN connections to each VPC for each of the company's development environment VRFs.

D. Configure a Transit Gateway Connect attachment for each MPLS VPN between the company's edge routers and Transit Gateway. Configure a transit gateway route table that matches the MPLS VPN for each of the company's development environments.

ANSWER: D

Explanation:

Configure a Transit Gateway Connect attachment for each MPLS VPN between the company's edge routers and Transit Gateway. Configure a transit gateway route table that matches the MPLS VPN for each of the company's development environments. This approach extends MPLS-style segmentation into the AWS Transit Gateway design by maintaining a distinct logical routing domain for each MPLS VPN. Transit Gateway Connect supports GRE and BGP-based connectivity, making it suitable for integrating routing domains from an MPLS environment while using the existing resilient Direct Connect connectivity. A separate Connect attachment and corresponding transit gateway route table for each development environment ensures that routes learned for one VPN remain isolated from routes in the other environments. Consequently, identical CIDR prefixes can exist in separate environments because they are not placed into, propagated through, or evaluated by the same routing table. As additional MPLS VPNs are created, the company can add the corresponding Connect attachment and route-table association without deploying, operating, and scaling a separate routing or SD-WAN appliance fleet. This uses Transit Gateway's managed route-table segmentation model and BGP route exchange to provide a scalable, low-administration mapping between on-premises VPN/VRF segmentation and AWS routing domains. See the AWS documentation for Transit Gateway Connect and Transit Gateway route tables.

QUESTION NO: 26

A company's security guidelines state that all outbound traffic from a VPC to the company's on-premises data center must pass through a security appliance. The security appliance runs on an Amazon EC2 instance. A network engineer needs to improve the network performance between the on-premises data center and the security appliance.

Which actions should the network engineer take to meet these requirements? (Choose two.)

A. Use an EC2 instance that supports enhanced networking.

B. Send outbound traffic through a transit gateway.

C. Increase the EC2 instance size.

D. Place the EC2 instance in a placement group within the VPC.

E. Attach multiple elastic network interfaces to the EC2 instance.

ANSWER: A C

Explanation:

Use an EC2 instance that supports enhanced networking and **Increase the EC2 instance size** directly improve the network capacity available to the EC2-based security appliance. Enhanced networking uses supported high-performance network interfaces and capabilities, including single-root I/O virtualization, to provide higher packets-per-second performance, lower latency, and reduced CPU overhead compared with legacy networking. This is particularly important for a security appliance, which must inspect and forward traffic without becoming a throughput bottleneck.

Increasing the instance size can also increase the instance's available network bandwidth, packet-processing capacity, and compute resources. EC2 instance network performance is determined in part by the instance type and size; larger sizes in an instance family generally provide greater networking capabilities. The appliance can therefore process more encrypted, inspected, or routed traffic between the VPC and the on-premises environment while maintaining its required traffic path. The selected instance type must support the appliance software and provide sufficient network performance for the expected workload. See the AWS documentation for enhanced networking on EC2 instances and EC2 instance network bandwidth.

QUESTION NO: 27

An international company wants to implement a multi-site hybrid infrastructure. The company wants to deploy its cloud computing resources on AWS in the us-east-1 Region and in the eu-west-2 Region, and in on-premises data centers in the United States (US) and in the United Kingdom (UK). The data centers are connected to each other by a private WAN connection. IP routing information is exchanged dynamically through BGP. The company wants to have two AWS Direct Connect connections, one each in the US and the UK.

The company expects to have 15 VPCs in each Region with CIDR blocks that do not overlap with each other or with CIDR blocks of the on-premises environment. The VPC CIDR blocks are planned so that the prefix aggregation can be performed both on a Regional level and across the entire AWS environment. The company will deploy a transit gateway in each Region to connect the VPCs. A network engineer plans to use a Direct Connect gateway in each Region. A transit VIF will attach the Direct Connect gateway in each Region to the transit gateway in that Region. The transit gateways will be peered with each other.

The network engineer wants to ensure that traffic follows the shortest geographical path from source to destination. Traffic between the on-premises data centers and AWS must travel across a local Direct Connect connection. Traffic between the US data center and eu-west-2 and traffic between the UK data center and us-east-1 must use the private WAN connection to reach the Direct Connect connection to the appropriate Region when the Direct Connect connection is available. The network must be resilient to failures in either the private WAN connection or with the Direct Connect connections. The network also must reroute traffic automatically in the event of any failure.

How should the network engineer configure the transit VIF associations on the Direct Connect gateways to meet these requirements?

A. Advertise only the aggregate route for the company's entire AWS environment.

B. Advertise VPC-specific CIDR prefixes from only the local Region. Additionally, advertise the aggregate route for the company's entire AWS environment.

C. Advertise all the specific VPC CIDR blocks from both Regions.

D. Advertise both Regional aggregate prefixes. Configure custom BGP communities on the routes advertised toward the data center.

ANSWER: B

Explanation:

Advertise VPC-specific CIDR prefixes from only the local Region. Additionally, advertise the aggregate route for the company's entire AWS environment. This design uses BGP longest-prefix matching to steer traffic geographically while retaining an automatic fallback path. At the US Direct Connect gateway, the specific us-east-1 VPC prefixes are advertised, along with a summarized prefix that represents all AWS VPC networks. The UK Direct Connect gateway similarly advertises the specific eu-west-2 prefixes and the overall AWS aggregate. Therefore, traffic for a local Regional VPC matches the more-specific route learned through the local Direct Connect connection. Traffic from the US data center to an eu-west-2 VPC matches the more-specific eu-west-2 route learned through the UK path across the private WAN; the reverse applies for UK-to-us-east-1 traffic. The AWS-wide aggregate remains available as a less-specific backup route. If the private WAN or the preferred Direct Connect path fails, BGP withdraws the affected more-specific routes and routing converges automatically to an available aggregate path, including a path through transit gateway peering where applicable. Direct Connect gateway associations support specifying allowed prefixes, and AWS documents the route propagation behavior for transit virtual interfaces and transit gateways. See AWS Direct Connect transit gateway associations and AWS Transit Gateway peering.

QUESTION NO: 28

A company deploys a software solution on Amazon EC2 instances that are in a clusterplacement group. The solution's UI is a single HTML page. The HTML file size is 1,024 bytes. The software processes files that exceed 1,024 MB in size. The software shares files over the network to clients upon request. The files are shared with the Don't Fragment flag set. Elastic network interfaces of the EC2 instances are set up with jumbo frames.

The UI is always accessible from all allowed source IP addresses, regardless of whether the source IP addresses are within a VPC, on the internet, or on premises. However, clients sometimes do not receive files that they request because the files fail to travel successfully from the software to the clients.

Which options provide a possible root cause of these failures? (Choose two.)

A. The source IP addresses are from on-premises hosts that are routed over AWS Direct Connect.

B. The source IP addresses are from on-premises hosts that are routed over AWS Site-to-Site VPN.

C. The source IP addresses are from hosts that connect over the public internet.

D. The security group of the EC2 instances does not allow ICMP traffic.

E. The operating system of the EC2 instances does not support jumbo frames.

ANSWER: B C D

Explanation:

The source IP addresses are from on-premises hosts that are routed over AWS Site-to-Site VPN, the source IP addresses are from hosts that connect over the public internet, and the security group of the EC2 instances does not allow ICMP traffic are all possible causes. Although the ENIs and cluster placement group can use jumbo frames within the applicable VPC path, the end-to-end path to a VPN customer gateway or through an internet gateway has a lower MTU than a jumbo-frame path. AWS documents a 1,500-byte MTU for traffic through an internet gateway and for Site-to-Site VPN connections. When the application sends packets with the Don't Fragment bit set, an intermediate device cannot fragment an oversized packet to fit such a path and must discard it instead.

Path MTU Discovery is intended to prevent this condition by returning ICMP “fragmentation needed” information to the sender, allowing the sender to reduce packet size. If the instance security group blocks the required inbound ICMP messages, the instance might not learn the lower path MTU. This can create a PMTUD black hole: small UI requests continue to succeed, while large file-transfer traffic fails. AWS recommends allowing the applicable ICMP traffic to support PMTUD. See Amazon EC2 network MTU documentation and AWS guidance for Path MTU Discovery.

QUESTION NO: 29

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 connections. Reroute traffic through the new public VIF.

B. 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.

C. Deploy a new pair of 10 GB Direct Connect connections with MACsec. Configure MACsec on the edge routers. Reroute traffic to the new Direct Connect connections. Decommission the original Direct Connect connections

D. Deploy a new pair of 10 GB Direct Connect connections with MACsec. Deploy a new public VIF on the new Direct Connect connections. Deploy two AWS Site-to-Site VPN connections on top of the new public VIF. Reroute traffic from the existing private VIF to the new Site-to-Site connections. Decommission the original Direct Connect connections.

ANSWER: C

Explanation:

Deploy a new pair of 10 GB Direct Connect connections with MACsec. Configure MACsec on the edge routers. Reroute traffic to the new Direct Connect connections. Decommission the original Direct Connect connections is correct because MACsec provides Layer 2 encryption for traffic traversing the physical Direct Connect links between the customer edge routers and AWS Direct Connect locations. This directly protects the previously unencrypted segment from the colocation edge routers to AWS, while preserving the high throughput required by the workload. AWS supports MACsec on dedicated 10 Gbps and 100 Gbps Direct Connect connections, so two new 10 Gbps connections can retain both the company’s existing aggregate bandwidth and its redundant connectivity design. MACsec must be configured on both the AWS Direct Connect connection and the customer’s compatible edge device. Creating replacement connections permits MACsec to be enabled and validated before production traffic is moved, minimizing disruption. After traffic has been rerouted to the protected redundant links, the original non-MACsec connections can be retired. This approach uses native link-layer encryption rather than encapsulating traffic in VPN tunnels, avoiding the operational management and potential throughput constraints associated with maintaining VPN encryption at this bandwidth. See the AWS Direct Connect MACsec documentation and the AWS Direct Connect User Guide.

QUESTION NO: 30

A company has two data centers that are interconnected with multiple redundant linksfrom different suppliers. The company uses IP addresses that are within the 172.16.0.0/16 CIDR block. The company is running iBGP between the two data centers by using a private Autonomous System Number (ASN)and IGP.

The company is moving toward a hybrid setup in which the company will initially use one VPC in the AWS Cloud. An AWS Direct Connect connection runs from the first data center to a Direct Connect gateway by using a private VIF On the connection, the company advertises a summarized route for the 172.16.0.0/16 network The company is planning to set up a second summarized route from the second data center to a different Direct Connect location.

The company needs to implement a solution to route traffic to and from AWS through the first Direct Connect connection. The solution must use the second Direct Connect connection for failover purposes only.

Which solution will meet these requirements?

A. Prepend the private ASN on the BGP announcements to AWS from the second data center. Add a second VIF in the first Direct Connect connection. Advertise the same network without any prepends from the first data center. Implement the same setup for the BGP announcement from AWS to the two data centers.

B. Tag the BGP announcements with the local preference BGP community tags. Set the tag to high preference for the first data center. Set the tag to low preference for the second data center. Configure the second data center's router to have a lower local preference for the direct AWS BGP advertisements than for the advertisement from the first data center.

C. Configure the Direct Connect gateway to prefer routing through the Direct Connect connection with the first data center. Configure the second data center's router to have a lower local preference for the direct AWS BGP advertisements than for the advertisement from the first data center.

D. Configure the local AWS Region BGP community tag on the BGP route that is advertised from the first data center. Configure AS PATH prepends on the BGP announcements from the second data center.

E. Prepend the private ASN to the 172.16.0.0/16 BGP advertisement from the second data center to AWS. On the second data center router, assign a lower local preference to AWS routes learned over its local Direct Connect connection than to the corresponding routes learned from the first data center.

ANSWER: E

Explanation:

Prepending the private ASN on the summarized 172.16.0.0/16 advertisement from the second data center makes the AS_PATH for that route longer when AWS compares the two otherwise equivalent routes. Consequently, AWS selects the route advertised through the first Direct Connect connection for traffic destined to the company network. If the first connection or its BGP route fails, the remaining route through the second location is selected automatically, providing the required active/passive behavior for AWS-to-on-premises traffic.

The complementary local-preference policy at the second data center controls the opposite traffic direction. The router there assigns a lower local preference to AWS routes learned directly through its own Direct Connect connection than to equivalent AWS routes received through the first data center. It therefore forwards on-premises-to-AWS traffic over the inter-data-center network and out through the first connection during normal operation. If that preferred route is withdrawn, the locally learned Direct Connect route remains available and becomes the failover path. This uses standard BGP policy controls in the appropriate direction: AS_PATH influences AWS route selection, while local preference controls the company router's outbound path selection. AWS describes Direct Connect BGP routing behavior in the Direct Connect routing and BGP documentation and recommends BGP-based resiliency designs in the Direct Connect Resiliency Toolkit.

QUESTION NO: 31

A company has a data center in the us-west-1 Region with a 10 Gbps AWS Direct Connect dedicated connection to a Direct Connect gateway. There are two private VIFs from the same data center location in us-west-1 that are attached to the same Direct Connect gateway.

VIF 1 advertises 172.16.0.0/16 with an AS PATH attribute value of 65000. VIF 2 advertises 172.16.1.0/24 with an AS PATH attribute value of 65000 65000 65000.

How will AWS route traffic to the data center for traffic that has a destination address within the 172.16.1.0/24 network range?

A. AWS will route all traffic by using VIF 1.

B. AWS will route all traffic by using VIF 2.

C. AWS will use both VIFs for routing by using a round-robin policy.

D. AWS will use flow control to balance the traffic between the two VIFs.

ANSWER: B

Explanation:

AWS will route all traffic by using VIF 2. For traffic whose destination is within 172.16.1.0/24, the route advertised through VIF 2 is a more specific prefix than the 172.16.0.0/16 route advertised through VIF 1. Route selection first uses the longest-prefix-match principle: a /24 route is preferred for addresses in that /24 over a covering /16 route. This decision applies even though the /24 advertisement has a longer AS_PATH due to repeated instances of ASN 65000.

The AS_PATH attribute is a BGP path-selection attribute and can influence selection when competing routes advertise the same destination prefix length. It does not override the fundamental IP forwarding lookup that selects the most specific matching route. Therefore, all traffic addressed to 172.16.1.0 through 172.16.1.255 is forwarded toward the data center over the private VIF that advertises 172.16.1.0/24. AWS documents Direct Connect route evaluation and BGP behavior in its Direct Connect routing and BGP documentation; its VPC route-table documentation also describes the longest-prefix-match routing principle.

QUESTION NO: 32

A company has an AWS Site-to-Site VPN connection between its office and its VPC. Users report occasional failure of the connection to the application that is hosted inside the VPC. A network engineer discovers in the customer gateway logs that the Internet Key Exchange (IKE) session ends when the connection to the application fails.

What should the network engineer do to bring up the IKE session if the IKE session goes down?

A. Set the dead peer detection (DPD) timeout action to Clear. Initiate traffic from the VPC to on premises.

B. Set the dead peer detection (DPD) timeout action to Restart. Initiate traffic from on premises to the VPC.

C. Set the dead peer detection (DPD) timeout action to None. Initiate traffic from the VPC to on premises.

D. Set the dead peer detection (DPD) timeout action to Cancel. Initiate traffic from on premises to the VPC.

ANSWER: B

Explanation:

Set the dead peer detection (DPD) timeout action to Restart. Initiate traffic from on premises to the VPC. DPD monitors the liveness of the peer for an AWS Site-to-Site VPN tunnel. Configuring the timeout action as *Restart* directs the AWS VPN endpoint to restart IKE when DPD determines that the peer is no longer responding. This makes the tunnel recovery behavior proactive rather than leaving the tunnel inactive after the IKE security association has ended.

Generating traffic from the customer gateway side toward a VPC resource supplies traffic that requires the IPsec tunnel and helps trigger IKE/IPsec negotiation on the customer gateway. Once the IKE security association is recreated, the tunnel can establish the required IPsec security associations and resume protected application traffic. The DPD timeout action is a per-tunnel Site-to-Site VPN option and should be configured consistently with the capabilities and expected recovery behavior of the customer gateway device. AWS documents the available DPD actions and their tunnel behavior in the Site-to-Site VPN tunnel options documentation. Additional operational guidance for diagnosing and restoring VPN tunnel connectivity is available in the AWS Knowledge Center.

QUESTION NO: 33

A company has developed an application on AWS that will track inventory levels of vending machines and initiate the restocking process automatically. The company plans to integrate this application with vending machines and deploy the vending machines in several markets around the world. The application resides in a VPC in the us-east-1 Region. The application consists of an Amazon Elastic Container Service (Amazon ECS) cluster behind an Application Load Balancer (ALB). The communication from the vending machines to the application happens over HTTPS.

The company is planning to use an AWS Global Accelerator accelerator and configure static IP addresses of the accelerator in the vending machines for application endpoint access. The application must be accessible only through the accelerator and not through a direct connection over the internet to the ALB endpoint.

Which solution will meet these requirements?

A. Configure the ALB in a private subnet of the VPC. Attach an internet gateway without adding routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to allow inbound traffic only from AWS Global Accelerator IP address ranges on the ALB listener port.

B. Configure the ALB in a private subnet of the VPC. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.

C. Configure the ALB in a public subnet of the VPAttach an internet gateway. Add routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.

D. Configure the ALB in a private subnet of the VPC. Attach an internet gateway. Add routes in thesubnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.

ANSWER: A

Explanation:

Configure the ALB in a private subnet, attach an internet gateway to the VPC, and do not add an internet-gateway route to the ALB subnet route table. This design supports an internal ALB endpoint for AWS Global Accelerator while preventing the ALB from being directly reachable through a public internet route. Global Accelerator requires the VPC containing an internal Application Load Balancer endpoint to have an internet gateway attached, which indicates that the VPC accepts internet traffic; however, the load balancer does not require a public IP address or a subnet route to the internet gateway.

The ALB security group must permit inbound HTTPS only from the AWS Global Accelerator service IP ranges, rather than from arbitrary internet sources. This ensures that traffic reaching the ALB is delivered through Global Accelerator and that clients cannot bypass the accelerator to access the ALB directly. The vending machines use the accelerator's static anycast IP addresses, gaining the intended globally distributed entry point while the application remains privately addressed in the VPC. See AWS Global Accelerator secure VPC connections and How AWS Global Accelerator works.

QUESTION NO: 34

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 premises. Associate the VPN attachment with this transit gateway route table. Propagate all application VPC attachments to this transit gateway route table.

B. Configure a separate transit gateway route table for each application VPC. Associate each application VPC attachment with its respective transit gateway route table. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.

C. Configure a separate transit gateway route table for all application VPCs. Associate all application VPCs with this transit gateway route table. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.

D. Configure a separate transit gateway route table for the shared services VPC. Associate the shared services VPC attachment with this transit gateway route table. Propagate all application VPC attachments to this transit gateway route table.

E. Configure a separate transit gateway route table for on premises and the shared services VPC. Associate the VPN attachment and the shared services VPC attachment with this transit gateway route table. Propagate all application VPC attachments to this transit gateway route table.

ANSWER: C E

Explanation:

Transit Gateway routing is controlled by the route table associated with the source attachment. To minimize route tables while preserving application-VPC isolation, all application VPC attachments can share one route table. In that table, propagating only the shared services VPC attachment and the VPN attachment installs routes to those destinations. Because no application VPC attachment routes are propagated into this table, traffic originating in one application VPC has no Transit Gateway route to another application VPC.

Return-path routing requires a second route table. Associating both the VPN attachment and the shared services VPC attachment with that route table, then propagating every application VPC attachment into it, provides routes from on premises and shared services back to each application VPC. This produces bidirectional application-to-on-premises and application-to-shared-services connectivity while using only two Transit Gateway route tables, regardless of the number of application VPCs.

Association determines which route table evaluates traffic entering through an attachment, while propagation dynamically adds an attachment's routes to a route table. AWS documents these Transit Gateway routing behaviors in the Transit Gateway route tables guide and illustrates this segmentation pattern in the centralized router documentation.

QUESTION NO: 35

A retail company is migrating its on-premises application to the AWS Cloud. Currently, the company has two on-premises data center locations. One data center is on the east coast of the United States, and one data center is on the west coast.

Each data center hosts four database systems. The largest database system stores 500 GB of data. The data centers are interconnected by two 10 GbE circuits for data synchronization. Each data center has two separate 1 GbE upstream internet connections. The company plans to have eight total VPCs to service its multiple business units. Four VPCs will be in the us-east-1 Region, and four will be in the us-west-2 Region.

A network engineer needs to design a connectivity solution that allows VPC-to-VPC connectivity. The solution must also allow secure connections between the on-premises data centers and AWS during the migration process. The company expects spikes in traffic among the VPCs during database synchronization. The company wants to run the migration plan during one weekend and as soon as technically possible. The company also wants to minimize long-term operational and human resources costs.

Which combination of steps will meet these requirements? {Select TWO.)

A. Deploy a transit gateway in each Region, attach all local VPCs to it, and peer the transit gateways. Update the transit gateway and VPC route tables to allow any VPC to connect to any other VPC.

B. Configure VPC peering between all the VPCs. Update the VPC route tables to allow connectivity.

C. Provision two AWS Direct Connect connections from two Direct Connect locations that serve us-east-1 and us-west-2 to provide connectivity between the data centers and AWS.

D. Provision one transit gateway VPN attachment for each data center to build connectivity between the on-premises data centers and AWS VPCs.

E. Provision one AWS Site-to-Site VPN connection for each data center and for each VPC to build connectivity between the on-premises data centers and AWS VPCs.

ANSWER: A C

Explanation:

Deploying a transit gateway in each AWS Region, attaching the local VPCs, and establishing inter-Region transit gateway peering provides a scalable hub-and-spoke network. Each transit gateway supports VPC attachments within its own Region, while transit gateway peering provides private routed connectivity between the Regional transit gateways. This removes the need to manage a large mesh of individual VPC connections and route-table relationships. Transit Gateway is designed to centralize network connectivity and can accommodate the expected bursts of inter-VPC database-synchronization traffic through its distributed architecture. AWS documents the Regional attachment model and inter-Region peering workflow in the Transit Gateway peering documentation.

Two AWS Direct Connect connections, using separate Direct Connect locations serving the two target Regions, provide dedicated private connectivity from the data centers to AWS. Direct Connect supports substantially greater and more predictable throughput than the existing 1 GbE internet links, enabling the migration data transfer to complete as quickly as the provisioned circuit capacity permits. The dedicated connections can be integrated with the transit-gateway-based design by using Direct Connect gateways and associated transit gateways. This approach also provides durable hybrid connectivity after the migration while reducing the operational burden of maintaining many point-to-point VPN connections. AWS describes Direct Connect connection capabilities, resiliency considerations, and available port speeds in the AWS Direct Connect User Guide.

QUESTION NO: 36

A network engineer is designing a hybrid networking environment that will connect a company's corporate network to the company's AWS environment. The AWS environment consists of 30 VPCs in 3 AWS Regions.

The network engineer needs to implement a solution to centrally filter traffic by using a firewall that the company's security team has approved. The solution must give all the VPCs the ability to connect to each other. Connectivity between AWS and the corporate network must meet a minimum bandwidth requirement of 2 Gbps.

Which solution will meet these requirements?

A. Deploy an IPsec VPN connection between the corporate network and a new transit gateway. Connect all VPCs to the transit gateway. Associate the approved firewall with the transit gateway.

B. Deploy a single 10 Gbps AWS Direct Connect connection between the corporate network and virtual private gateway of each VPC. Connect the virtual private gateways to a Direct Connect gateway. Build an IPsec tunnel to a new transit VPC. Deploy the approvedfirewall to the transit VPC.

C. Deploy two 1 Gbps AWS Direct Connect connections in different Direct Connect locations to connect to the corporate network. Build a transit VIF on each connection to a Direct Connect gateway. Associate the Direct Connect gateway with a new transit gateway for each Region. Configure the VIFs to use equal-cost multipath (ECMP) routing. Connect all the VPCs in the three Regions to the transit gateway. Configure the transit gateway route table

D. Deploy four 1 Gbps AWS Direct Connect connections in different Direct Connect locations to connect to the corporate network. Build a transit VIF on each connection to a Direct Connect gateway. Associate the Direct Connect gateway with a new transit gateway for each Region. Connect the transit gateways by using a transit gateway peering attachment. Configure the VIFs to use equal-cost multipath (ECMP) routing. Configure transit gateway route

ANSWER: D

Explanation:

Deploying four 1 Gbps AWS Direct Connect connections across different Direct Connect locations, using transit virtual interfaces to a Direct Connect gateway, provides resilient hybrid connectivity and sufficient aggregate bandwidth. The Direct Connect gateway can associate with transit gateways in each required Region, allowing the corporate network to reach workloads attached to those transit gateways. Transit gateway peering then provides inter-Region connectivity so that all 30 VPCs can communicate across the three Regions.

Equal-cost multipath routing across the Direct Connect transit VIFs enables traffic to use multiple available paths. With four 1 Gbps connections, the design has at least 2 Gbps of aggregate capacity while also providing connection and location redundancy. A centrally deployed, approved firewall can be inserted into the traffic path through transit gateway route-table configuration. Routes can direct VPC-to-VPC and hybrid traffic through firewall inspection attachments before forwarding

traffic toward destination VPCs, peer transit gateways, or the Direct Connect gateway. This combines scalable VPC connectivity, centralized inspection, multi-Region routing, and the required hybrid bandwidth. See AWS Direct Connect gateway and transit gateway associations and centralized routing with AWS Transit Gateway.

QUESTION NO: 37

A company has agreed to collaborate with a partner for a research project. The company has multiple VPCs in the us-east-1 Region that use CIDR blocks within 10.10.0.0/16. The VPCs are connected by a transit gateway that is named TGW-C in us-east-1. TGW-C has an Autonomous System Number (ASN) configuration value of 64520.

The partner has multiple VPCs in us-east-1 that use CIDR blocks within 172.16.0.0/16. The VPCs are connected by a transit gateway that is named TGW-P in us-east-1. TGW-P has an ASN configuration value of 64530.

A network engineer needs to establish network connectivity between the company's VPCs and the partner's VPCs in us-east-1.

Which solution will meet these requirements with MINIMUM changes to both networks?

A. Create a new VPC in a new account. Deploy a router from AWS Marketplace. Share TGW-C and TGW-P with the new account by using AWS Resource Access Manager (AWS RAM). Associate TGW-C and TGW-P with the new VPC.

B. Create an IPsec VPN connection between TGW-C and TGW-P. Configure the routing between the transit gateways to use the IPsec VPN connection.

C. Configure a cross-account transit gateway peering attachment between TGW-C and TGW-P. Configure the routing between the transit gateways to use the peering attachment.

D. Share TGW-C with the partner account by using AWS Resource Access Manager (AWS RAM). Associate the partner VPCs with TGW-C.

ANSWER: C

Explanation:

Configure a cross-account transit gateway peering attachment between TGW-C and TGW-P. Configure the routing between the transit gateways to use the peering attachment. This directly connects the two existing transit gateway routing domains while allowing each organization to retain its current VPC attachments, CIDR ranges, and transit gateway design. AWS Transit Gateway supports peering attachments between transit gateways in different AWS accounts and in the same Region. The owner of one transit gateway creates the peering attachment, and the owner of the peer transit gateway accepts it. After acceptance, each organization adds routes in the applicable transit gateway route tables for the remote network prefixes through the peering attachment. Here, routes for the partner's 172.16.0.0/16 address space can be directed from the company transit gateway to the peer attachment, and routes for the company's 10.10.0.0/16 address space can be directed in the reverse direction. The distinct, non-overlapping address ranges permit unambiguous routing between the networks. This approach is purpose-built for connecting independently administered transit gateways and requires only the peering attachment, acceptance, and route-table updates, minimizing changes to both existing networks. See the AWS Transit Gateway peering guide and AWS Transit Gateway route table documentation.

QUESTION NO: 38

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

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

A. Use an Amazon Route 53 Resolver inbound endpoint.

B. Modify the DHCP options set by setting a custom DNS server value.

C. Use an Amazon Route 53 Resolver outbound endpoint.

D. Create DNS proxy servers.

E. Create Amazon Route 53 private hosted zones.

F. Set up a zone transfer between Amazon Route 53 and the on-premises DNS.

ANSWER: A C E

Explanation:

Use an Amazon Route 53 Resolver inbound endpoint, **use an Amazon Route 53 Resolver outbound endpoint**, and **create Amazon Route 53 private hosted zones**. Private hosted zones provide AWS-native, VPC-scoped authoritative DNS namespaces, allowing application teams to manage records for their development applications without operating DNS servers. They can use separate delegated-style application namespaces or distinct private zones as required by the organization's naming model.

A Route 53 Resolver inbound endpoint enables DNS resolvers in the on-premises network to submit queries to Route 53 Resolver. This gives on-premises clients resolution of records hosted in the AWS private hosted zones. A Route 53 Resolver outbound endpoint provides the reverse direction: it sends queries from VPC workloads to on-premises DNS servers. Resolver rules associated with the relevant VPCs direct queries for centrally managed on-premises domains to that outbound endpoint. Together, the endpoints and rules provide the required bidirectional hybrid DNS resolution while retaining managed Route 53 Resolver infrastructure rather than requiring customer-managed DNS proxy fleets.

For implementation details, see the Amazon Route 53 Resolver documentation and Route 53 private hosted zones documentation.

QUESTION NO: 39

A company has a hybrid cloud environment. The company's data center is connected to the AWS Cloud by an AWS Direct Connect connection. The AWS environment includes VPCs that are connected together in a hub-and-spoke model by a transit gateway. The AWS environment has a transit VIF with a Direct Connect gateway for on-premises connectivity.

The company has a hybrid DNS model. The company has configured Amazon Route 53 Resolver endpoints in the hub VPC to allow bidirectional DNS traffic flow. The company is running a backend application in one of the VPCs.

The company uses a message-oriented architecture and employs Amazon Simple Queue Service (Amazon SQS) to receive messages from other applications over a private network. A network engineer wants to use an interface VPC endpoint for Amazon SQS for this architecture. Client services must be able to access the endpoint service from on premises and from multiple VPCs within the company's AWS infrastructure.

Which combination of steps should the network engineer take to ensure that the client applications can resolve DNS for the interface endpoint? (Choose three.)

A. Create the interface endpoint for Amazon SQS with the option for private DNS names turned on.

B. Create the interface endpoint for Amazon SQS with the option for private DNS names turned off.

C. Manually create a private hosted zone for sqs.us-east-1.amazonaws.com. Add necessary records that point to the interface endpoint. Associate the private hosted zones with other VPCs.

D. Use the automatically created private hosted zone for sqs.us-east-1.amazonaws.com with previously created necessary records that point to the interface endpoint. Associate the private hosted zones with other VPCs.

E. Access the SQS endpoint by using the public DNS name sqs.us-east-1.amazonaws.com in VPCs and on premises.

F. Access the SQS endpoint by using the private DNS name of the interface endpoint .sqs.us-east-1.vpce.amazonaws.com in VPCs and on premises.

ANSWER: B C E

Explanation:

The interface endpoint for Amazon SQS should be created with private DNS disabled because the endpoint is centralized in the hub VPC but must be consumed from multiple spoke VPCs and from on premises. The AWS-managed private hosted zone that is created when private DNS is enabled is associated only with the endpoint VPC; instead, the company needs a Route 53 private hosted zone that it controls and can associate with all required VPCs.

Create a private hosted zone for `sqs.us-east-1.amazonaws.com` and create the required alias record targeting the interface endpoint's regional DNS name. Associate this private hosted zone with the hub and spoke VPCs. Queries for the normal regional SQS hostname will then resolve to the endpoint's private IP addresses for workloads in those VPCs. The existing Route 53 Resolver endpoints provide the hybrid DNS path so on-premises clients can send queries through the appropriate Resolver endpoint and receive the same private resolution. Applications should therefore use the standard SQS regional hostname, `sqs.us-east-1.amazonaws.com`, rather than an endpoint-specific hostname. AWS documents this centralized PrivateLink DNS pattern in its Transit Gateway, PrivateLink, and Route 53 Resolver guidance and describes interface endpoint DNS behavior in the AWS PrivateLink documentation.

QUESTION NO: 40

A network engineer is designing hybrid connectivity with AWS Direct Connect and AWS Transit Gateway. A transit gateway is attached to a Direct Connect gateway and 19 VPCs across different AWS accounts. Two new VPCs are being attached to the transit gateway. The IP address administrator has assigned 10.0.32.0/21 to the first VPC and 10.0.40.0/21 to the second VPC. The prefix list has one CIDR block remaining before the prefix list reaches the quota for the maximum number of entries.

What should the network engineer do to advertise the routes from AWS to on premises to meet these requirements?

A. Add 10.0.32.0/21 and 10.0.40.0/21 to both AWS managed prefix lists.

B. Add 10.0.32.0/21 and 10.0.40.0/21 to the allowed prefix list.

C. Add 10.0.32.0/20 to both AWS managed prefix lists.

D. Add 10.0.32.0/20 to the allowed prefix list.

ANSWER: D

Explanation:

Add 10.0.32.0/20 to the allowed prefix list. The two assigned VPC ranges are contiguous: 10.0.32.0/21 covers addresses from 10.0.32.0 through 10.0.39.255, and 10.0.40.0/21 covers 10.0.40.0 through 10.0.47.255. These two adjacent networks can therefore be summarized into the single supernet 10.0.32.0/20, which covers the complete address range from 10.0.32.0 through 10.0.47.255. This consumes only one remaining allowed-prefix entry while ensuring that both VPC CIDR ranges are included in the routes advertised through the Direct Connect gateway to the on-premises network.

For a transit gateway association with a Direct Connect gateway, allowed prefixes define the CIDR blocks that the transit gateway can advertise to the Direct Connect gateway. Route summarization is appropriate here because the aggregate is correctly aligned on a /20 boundary and does not include address space outside the two newly assigned, contiguous VPC ranges. AWS documents the allowed-prefix behavior for Direct Connect gateway transit gateway associations at AWS Direct Connect: Associate a transit gateway with a Direct Connect gateway and CIDR routing fundamentals at Amazon VPC subnet and CIDR block sizing.

QUESTION NO: 41

A company's AWS infrastructure is spread across more than 50 accounts and across five AWS Regions. The company needs to manage its security posture with simplified administration and maintenance for all the AWS accounts. The company wants to use AWS Firewall Manager to manage the firewall rules and requirements.

The company creates an organization with all features enabled in AWS Organizations.

Which combination of steps should the company take next to meet the requirements? (Select THREE.)

A. Configure only the Firewall Manager administrator account to join the organization.

B. Configure all the accounts to join the organization.

C. Set an account as the Firewall Manager administrator account.

D. Set an account as the Firewall Manager child account.

E. Set up AWS Config for all the accounts and all the Regions where the company has resources.

F. Set up AWS Config for only the organization's management account.

ANSWER: B C E

Explanation:

AWS Firewall Manager is designed to centrally deploy and maintain security policies across accounts that belong to an AWS Organization. Therefore, all workload accounts that require centrally managed firewall policies must be members of the organization. This enables Firewall Manager policies to be applied consistently across the organization or to selected organizational units and accounts.

The company must also designate a Firewall Manager administrator account. AWS Organizations uses delegated administration for this purpose: the organization's management account registers a member account as the delegated administrator, and that account then creates and administers Firewall Manager policies for the organization. This separates centralized security administration from day-to-day management of the organization.

Finally, AWS Config must be enabled in every in-scope account and AWS Region where the company has resources. Firewall Manager relies on AWS Config to discover resources, evaluate policy compliance, and identify resources that require remediation. Enabling Config broadly across the organization and all relevant Regions provides the resource inventory and compliance visibility needed for centralized policy enforcement. See the AWS Firewall Manager prerequisites documentation and the AWS Organizations delegated administrator documentation.

QUESTION NO: 42

A company has started using AWS Cloud WAN with one edge location in the us-east-1 Region. The company has a production segment and a security segment in AWS Cloud WAN. The company also has a default core network policy.

The company has created a production VPC for the production workload. The company has created an outbound inspection VPC to inspect internet-bound traffic from the production VP

C.

The company has attached the production VPC to the production segment and has attached the outbound inspection VPC to the security segment. The company has also created an AWS Network Firewall firewall in the outbound inspection VPC to inspect internet-based traffic.

The company has updated a route table for the production VPC to send all internet-bound traffic to the AWS Cloud WAN core network. The company has updated a route table for the outbound inspection VPC to ensure that Network Firewall inspects any outgoing traffic and incoming traffic.

During testing, an Amazon EC2 instance in the production VPC cannot reach the internet. The company checks the Network Firewall rules and confirms that the rules are not blocking the traffic.

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

A. Update the core network policy to configure segment sharing. Share the production segment with the security segment.

B. Update the core network policy to create a static route for the security segment. Specify 0.0.0.0/0 as the destination CIDR block. Specify the outbound inspection VPC as an attachment.

C. Update the core network policy to create a static route for the production segment. Specify 0.0.0.0/0 as the destination CIDR block. Specify the outbound inspection VPC as an attachment.

D. Update the core network policy to create a static route for the production segment.Specify 10.2.0.0/16 as the destination CIDR block. Specify the outbound inspection VPC as an attachment.

E. Create an attachment to attach the outbound inspection VPC to the production segment. Update the core network policy to turn on isolated attachment for the production segment.

ANSWER: A C

Explanation:

The production segment must be able to use the attachment that is associated with the security segment, because the outbound inspection VPC is the intended centralized egress point. Configuring segment sharing from the production segment to the security segment allows the production segment to use routes and attachments made available by the security segment. This establishes the required cross-segment connectivity model in the Cloud WAN core network policy.

The production segment also needs an explicit default route for `0.0.0.0/0` with the outbound inspection VPC attachment as its destination. The production VPC already directs internet-bound traffic to Cloud WAN, but Cloud WAN requires a route that selects the inspection VPC attachment as the next hop for that traffic. The inspection VPC can then send the traffic through the Network Firewall endpoint according to its VPC route tables and onward to its internet egress path. Together, segment sharing and the production-segment default route provide a controlled route from the workload VPC to the centralized inspection VPC while preserving firewall inspection for the egress flow.

AWS documents Cloud WAN segment sharing and static routes in the core network policy reference. For the required routing design inside the inspection VPC, see AWS Network Firewall internet gateway routing architectures.

QUESTION NO: 43

A company plans to run a computationally intensive data processing application on AWS. The data is highly sensitive. The VPC must have no direct internet access, and the company has applied strict network security to control access.

Data scientists will transfer data from the company's on-premises data center to the instances by using an AWS Site-to-Site VPN connection. The on-premises data center uses the network range 172.31.0.0/20 and will use the network range 172.31.16.0/20 in the application VPC.

The data scientists report that they can start new instances of the application but that they cannot transfer any data from the on-premises data center. A network engineer enables VPC flow logs and sends a ping to one of the instances to test reachability. The flow logs show the following:

The network engineer must recommend a solution that will give the data scientists the ability to transfer data from the on-premises data center.

Which solution will meet these requirements?

A. Modify the security group for the application. Add an inbound rule to allow traffic from the on-premises data center network range to the application.

B. Modify the network ACLs for the VPC subnet. Add an inbound rule to allow traffic from the on-premises data center network range to the VPC subnet range.

C. Modify the network ACLs for the VPC subnet. Add an outbound rule to allow traffic from the VPC subnet range to the on-premises data center network range.

D. Modify the security group for the application. Add an outbound rule to allow traffic from the application to the on-premises data center network range.

ANSWER: C

Explanation:

Modify the network ACLs for the VPC subnet. Add an outbound rule to allow traffic from the VPC subnet range to the on-premises data center network range. Network ACLs are stateless filters, so they evaluate inbound traffic and the corresponding return traffic independently. When a host in 172.31.0.0/20 initiates communication with an application instance in 172.31.16.0/20 through the Site-to-Site VPN, the instance must send response packets back to the on-premises address range. The response uses the VPC subnet as its source and the on-premises network as its destination. Therefore, the subnet network ACL requires an outbound rule that permits this return path, including the applicable protocol and

ephemeral port range for the data-transfer traffic. This enables bidirectional communication without creating direct internet access to the VPC. AWS VPC Flow Logs can show that packets are accepted in one direction but rejected by the network ACL in the reverse direction, which is consistent with a missing stateless return rule. AWS documents the stateless behavior and the need for rules in both directions for network ACLs in the Network ACLs guide. See also AWS guidance on Site-to-Site VPN connectivity.