

# ValidBraindumps

Over **61842+** Satisfied Customers

About Us

ValidBraindumps

HOME CERTIFICATIONS HOW TO PAY? GUARANTEE FAQ CART (0)

Test4engine

WE

TRY BEFORE YOU BUY

Download a free sample of any of our exam questions and answers

- > 24/7 customer support, Secure shopping site
- > Free One year updates to match real exam scenarios
- > If you failed your exam after buying our products we will refund the full amount back to you.

IGINE

st version

n exams. Besides for the

ear test questions and

to the highest standards of technical

fronter experts and published authors for

Select a vendor... Select an exam... Your email address Free Download

## Try Before You Buy

Download a free sample of any of our exam questions and answers

- > 24/7 customer support, Secure shopping site
- > Free One year updates to match real exam scenarios
- > If you failed your exam after buying our products we will refund the full amount back to you.

<http://www.validbraindumps.com>

Free valid test braindumps for IT certification valid exam

**Exam** : **ANS-C01**

**Title** : AWS Certified Advanced  
Networking Specialty Exam

**Vendor** : Amazon

**Version** : DEMO

**NO.1** 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

**NO.2** A company has hundreds of Amazon EC2 instances that are running in two production VPCs across all Availability Zones in the us-east-1 Region. The production VPCs are named VPC A and VPC B.

A new security regulation requires all traffic between production VPCs to be inspected before the traffic is routed to its final destination. The company deploys a new shared VPC that contains a stateful firewall appliance and a transit gateway with a VPC attachment across all VPCs to route traffic between VPC A and VPC B through the firewall appliance for inspection. During testing, the company notices that the transit gateway is dropping the traffic whenever the traffic is between two Availability Zones.

What should a network engineer do to fix this issue with the LEAST management overhead?

**A.** In the shared VPC, replace the VPC attachment with a VPN attachment. Create a VPN tunnel between the transit gateway and the firewall appliance. Configure BGP.

**B.** Enable transit gateway appliance mode on the VPC attachment in VPC A and VPC B.

- C. Enable transit gateway appliance mode on the VPC attachment in the shared VPC.
- D. In the shared VPC, configure one VPC peering connection to VPC A and another VPC peering connection to VPC B.

**Answer:** C

**NO.3** A network engineer configures a second AWS Direct Connect connection to an existing network. The network engineer runs a test in the AWS Direct Connect Resiliency Toolkit on the connections. The test produces a failure. During the failover event, the network engineer observes a 90-second interruption before traffic shifts to the failover connection.

Which solution will reduce the time for failover?

- A. Decrease the BGP hello timer to 5 seconds.
- B. Add a VPN connection to the connectivity solution. Implement fast failover.
- C. Configure Bidirectional Forwarding Detection (BFD) on the on-premises router.
- D. Decrease the BGP hold-down timer to 5 seconds.

**Answer:** C

**NO.4** 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. Configure a VPN gateway in each VPC. Create an AWS Site-to-Site VPN tunnel from each VPC to the transit VPC. Use BGP routing to route traffic between the VPCs and the transit VPC.
- D. Configure AWS Site-to-Site VPN connections between each VPC. Enable route propagation for each Site-to-Site VPN connection to route traffic between the VPCs.

**Answer:** B

**NO.5** 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 B E

**NO.6** A global company is establishing network connections between the company's primary and secondary data centers and a VPC. A network engineer needs to maximize resiliency and fault tolerance for the connections.

The network bandwidth must be greater than 10 Gbps.

Which solution will meet these requirements MOST cost-effectively?

- A.** Set up a 100 Gbps connection at the primary data center that terminates at an AWS Direct Connect location. Set up a second 100 Gbps connection at the secondary data center that terminates at a second Direct Connect location. Ensure the connections are managed by separate providers.
- B.** Set up a 10 Gbps connection at the primary data center that terminates at an AWS Direct Connect location. Set up a second 10 Gbps connection at the secondary data center that terminates at a second Direct Connect location. Ensure the connections are managed by separate providers.
- C.** Set up two 10 Gbps connections at the primary data center that terminate at one AWS Direct Connect location. Ensure the connections are managed by separate providers. Set up two 10 Gbps connections at the secondary data center that terminate at a second Direct Connect location. Ensure the connections are managed by separate providers.
- D.** Set up a 10 Gbps connection at the primary data center that terminates at an AWS Direct Connect location. Set up an AWS Site-to-Site VPN connection at the secondary data center that terminates at a virtual private gateway in the same Region as the company's VPC.

**Answer:** C

Explanation:

**Multiple 10 Gbps Connections:** By setting up two 10 Gbps connections at each data center, the solution achieves an aggregate bandwidth of 20 Gbps per data center, exceeding the requirement of 10 Gbps. Using multiple 10 Gbps links is more cost-effective than deploying 100 Gbps links.

**Separate Providers:** Ensuring that the connections are managed by separate providers minimizes the risk of a single provider's failure affecting the network.

**Two Direct Connect Locations:** Terminating connections at two Direct Connect locations ensures geographic redundancy. This setup minimizes the impact of outages or disruptions at a single Direct Connect location.

**NO.7** 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

**NO.8** A company has an application that runs on a fleet of Amazon EC2 instances. A new company regulation mandates that all network traffic to and from the EC2 instances must be sent to a centralized third-party EC2 appliance for content inspection.

Which solution will meet these requirements?

- A.** Configure VPC flow logs on each EC2 network interface. Publish the flow logs to an Amazon S3 bucket. Create a third-party EC2 appliance to acquire flow logs from the S3 bucket. Log in to the appliance to monitor network content.
- B.** Create a third-party EC2 appliance in an Auto Scaling group fronted by a Network Load Balancer (NLB). Configure a mirror session. Specify the NLB as the mirror target. Specify a mirror filter to capture inbound and outbound traffic for the source of the mirror session, specify the EC2 elastic network interfaces for all the instances that host the application.
- C.** Configure a mirror session. Specify an Amazon Data Firehose delivery stream as the mirror target. Specify a mirror filter to capture inbound and outbound traffic. For the source of the mirror session, specify the EC2 elastic network interfaces for all the instances that host the application. Create a third-party EC2 appliance. Send all traffic to the appliance through the Firehose delivery stream for content inspection.
- D.** Configure VPC flow logs on each EC2 network interface. Send the logs to Amazon CloudWatch. Create a third-party EC2 appliance. Configure a CloudWatch filter to send the flow logs to Amazon Data Firehose to load the logs into the appliance.

**Answer:** D

**NO.9** A company is planning to host external websites on AWS. The websites will include multiple tiers such as web servers, application logic services, and databases. The company wants to use AWS Network Firewall.

AWS IAM and VPC security groups for network security.

The company must ensure that the Network Firewall firewalls are deployed appropriately within relevant VPCs. The company needs the ability to centrally manage policies that are deployed to Network Firewall and AWS WAF rules. The company also needs to allow application teams to manage their own security groups while ensuring that the security groups do not allow overly permissive access.

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

- A.** Define Network Firewall firewalls, AWS WAFv2 web ACLs, Network Firewall policies, and VPC security groups in code. Use AWS CloudFormation to deploy the objects and initial policies and rule groups. Use CloudFormation to update the AWS WAFv2 web ACLs, Network Firewall policies, and VPC security groups. Use Amazon GuardDuty to monitor for overly permissive rules.
- B.** Define Network Firewall firewalls, AWS WAFv2 web ACLs, Network Firewall policies, and VPC security groups in code. Use the AWS Management Console or the AWS CLI to manage the AWS WAFv2 web ACLs, Network Firewall policies, and VPC security groups. Use Amazon GuardDuty to invoke an AWS Lambda function to evaluate the configured rules and remove any overly permissive rules.
- C.** Deploy AWS WAFv2 IP sets and AWS WAFv2 web ACLs with AWS CloudFormation. Use AWS Firewall Manager to deploy Network Firewall firewalls and VPC security groups where required and to manage the AWS WAFv2 web ACLs, Network Firewall policies, and VPC security groups.

**D.** Define Network Firewall firewalls. AWS WAFv2 web ACLs. Network Firewall policies, and VPC security groups in code. Use AWS CloudFormation to deploy the objects and initial policies and rule groups. Use AWS Firewall Manager to manage the AWS WAFv2 web ACLs. Network Firewall policies, and VPC security groups. Use Amazon GuardDuty to monitor for overly permissive rules.

**Answer:** D

**NO.10** A company is planning to use Amazon S3 to archive financial data. The data is currently stored in an on- premises data center. The company uses AWS Direct Connect with a Direct Connect gateway and a transit gateway to connect to the on-premises data center. The data cannot be transported over the public internet and must be encrypted in transit.

Which solution will meet these requirements?

**A.** Create a Direct Connect public VIF. Set up an IPsec VPN connection over the public VIF to access Amazon S3. Use HTTPS for communication.

**B.** Create an IPsec VPN connection over the transit VIF. Create a VPC and attach the VPC to the transit gateway. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.

**C.** Create a VPC and attach the VPC to the transit gateway. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.

**D.** Create a Direct Connect public VIF. Set up an IPsec VPN connection over the public VIF to the transit gateway. Create an attachment for Amazon S3. Use HTTPS for communication.

**Answer:** B

Explanation:

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

An IPsec VPN connection over the transit VIF can encrypt traffic between the on-premises network and AWS without using public IP addresses or the internet2. A VPC endpoint for Amazon S3 can enable private access to S3 buckets within the same region. HTTPS can provide additional encryption for communication.

**NO.11** An application team for a startup company is deploying a new multi-tier application into the AWS Cloud. The application will be hosted on a fleet of Amazon EC2 instances that run in an Auto Scaling group behind a publicly accessible Network Load Balancer (NLB). The application requires the clients to work with UDP traffic and TCP traffic.

In the near term, the application will serve only users within the same geographic location. The application team plans to extend the application to a global audience and will move the deployment to multiple AWS Regions around the world to bring the application closer to the end users. The application team wants to use the new Regions to deploy new versions of the application and wants to be able to control the amount of traffic that each Region receives during these rollouts. In addition, the application team must minimize first- byte latency and jitter (randomized delay) for the end users.

How should the application team design the network architecture for the application to meet these requirements?

**A.** Create an Amazon CloudFront distribution to align to each Regional deployment. Set the NLB for each Region as the origin for each CloudFront distribution. Use an Amazon Route 53 weighted routing policy to control traffic to the newer Regional deployments.

**B.** Create an AWS Global Accelerator accelerator and listeners for the required ports. Configure

endpoint groups for each Region. Configure a traffic dial for the endpoint groups to control traffic to the newer Regional deployments. Register the NLBs with the endpoint groups.

**C.** Use Amazon S3 Transfer Acceleration for the application in each Region. Adjust the amount of traffic that each Region receives from the Transfer Acceleration endpoints to the Regional NLBs.

**D.** Create an Amazon CloudFront distribution that includes an origin group. Set the NLB for each Region as the origins for the origin group. Use an Amazon Route 53 latency routing policy to control traffic to the new Regional deployments.

**Answer:** B

**NO.12** A company is in the early stage of AWS Cloud adoption. The company has an application that is running in an on-premises data center in Asia. The company needs to deploy new applications in the us-east-1 Region.

The applications in the cloud need connectivity to the on-premises data center.

The company needs to set up a communication channel between AWS and the data center. The solution must improve latency, minimize the possibility of performance impact from transcontinental routing over the public internet, and encrypt data in transit.

Which solution will meet these requirements in the LEAST amount of time?

**A.** Create an AWS Site-to-Site VPN connection with acceleration turned on. Create a virtual private gateway. Attach the Site-to-Site VPN connection to the virtual private gateway. Attach the virtual private gateway to the VPC where the applications will be deployed.

**B.** Create an AWS Site-to-Site VPN connection with acceleration turned on. Create a transit gateway. Attach the Site-to-Site VPN connection to the transit gateway. Create a transit gateway attachment to the VPC where the applications will be deployed.

**C.** Create an AWS Direct Connect connection. Create a virtual private gateway. Create a public VIF and a private VIF that use the virtual private gateway. Create an AWS Site-to-Site VPN connection over the public VIF.

**D.** Create an AWS Site-to-Site VPN connection with acceleration turned off. Create a transit gateway. Attach the Site-to-Site VPN connection to the transit gateway. Create a transit gateway attachment to the VPC where the applications will be deployed.

**Answer:** B

**NO.13** A company is developing an API-based application on AWS for its process workflow requirements. The API will be invoked by clients in the company's on-premises datacenters. The company has set up an AWS Direct Connect connection between on premises and AWS. A network engineer decides to implement the API as a private REST API in Amazon API Gateway. The network engineer wants to ensure that clients can reach the API endpoint through private communication. Which solution can the network engineer use to invoke the API without any additional infrastructure setup?

**A.** Create an interface VPC endpoint for API Gateway with private DNS names enabled. Access the API by using the private DNS name of the endpoint.

**B.** Create an interface VPC endpoint for API Gateway with private DNS names enabled. Access the API by using an Amazon Route 53 alias of the endpoint.

**C.** Create an interface VPC endpoint for API Gateway. Associate the endpoint with the private REST API.

Access the API by using an Amazon Route 53 alias of the endpoint.

**D.** Create an interface VPC endpoint for API Gateway with private DNS names enabled. Access the API by using the public DNS name of the endpoint.

**Answer:** A

**NO.14** A company runs applications in two VPCs that are in separate AWS Regions. One VPC is in the us-east-1 Region. The second VPC is in the us-west-1 Region. The company needs to establish connectivity between the two VPCs. The company also needs to connect the VPCs to applications that run in an on-premises data center.

The current traffic requirement between the VPCs is 50 TB per month. The company expects traffic volume between the VPCs to increase. The traffic requirement from the VPCs to the on-premises data center is 10 TB per month. The company expects the traffic between the VPCs and the data center to remain constant.

Which solution will meet these requirements MOST cost-effectively?

**A.** Create a transit gateway in each Region. Create VPN connections from the transit gateways to the on-premises firewall. Create a peering connection between the transit gateways.

**B.** Create a virtual private gateway in each Region. Create VPN connections from the on-premises firewall to the virtual private gateways. Configure the on-premises firewall to route the traffic between the two VPCs.

**C.** Create a virtual private gateway in each Region. Create VPN connections from the on-premises firewall to the virtual private gateways. Create a VPC peering connection between the two VPCs.

**D.** Create a virtual private gateway in each Region. Create VPN connections from the on-premises firewall to the virtual private gateways. Create a VPN connection between the virtual private gateways.

**Answer:** A

Explanation:

**Traffic Volume Consideration:** The traffic volume between the VPCs (50 TB per month and increasing) justifies the use of transit gateways, which are designed for scalable, high-throughput interconnectivity. A VPC peering connection would not scale as efficiently for this traffic volume.

**On-Premises Connectivity:** Establishing VPN connections from the on-premises firewall to the transit gateways ensures secure connectivity between the on-premises data center and both VPCs.

**Transit Gateway Peering:** Creating a peering connection between the transit gateways allows for efficient inter-Region communication between the VPCs without routing through the on-premises data center, reducing latency and costs.

**Cost Efficiency:** Transit gateway peering provides a cost-effective solution for large inter-Region traffic volumes compared to alternatives like routing all traffic through the on-premises data center, which would incur higher egress costs and potentially create a bottleneck.

**NO.15** A global film production company uses the AWS Cloud to encode and store its video content 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

**NO.16** A company uses Amazon Route 53 for its DNS needs. The company's security team wants to update the DNS infrastructure to provide the most recent security posture.

The security team has configured DNS Security Extensions (DNSSEC) for the domain. The security team wants a network engineer to explain who is responsible for the rotation of DNSSEC keys.

Which explanation should the network administrator provide to the security team?

- A. AWS rotates the zone-signing key (ZSK). The company rotates the key-signing key (KSK).
- B. The company rotates the zone-signing key (ZSK) and the key-signing key (KSK).
- C. AWS rotates the AWS Key Management Service (AWS KMS) key and the key-signing key (KSK).
- D. The company rotates the AWS Key Management Service (AWS KMS) key. AWS rotates the key-signing key (KSK).

**Answer:** A

**NO.17** Your company runs an application for the US market in the us-east-1 AWS region. This application uses proprietary TCP and UDP protocols on Amazon Elastic Compute Cloud (EC2) instances. End users run a real- time, front-end application on their local PCs. This front-end application knows the DNS hostname of the service.

You must prepare the system for global expansion. The end users must access the application with lowest latency.

How should you use AWS services to meet these requirements?

- A. Register the IP addresses of the service hosts as "A" records with latency-based routing policy in Amazon Route 53, and set a Route 53 health check for these hosts.
- B. Set the Elastic Load Balancing (ELB) load balancer in front of the hosts of the service, and register the ELB name of the main service host as an ALIAS record with a latency-based routing policy in Route 53.
- C. Set Amazon CloudFront in front of the host of the service, and register the CloudFront name of the main service as an ALIAS record in Route 53.
- D. Set the Amazon API gateway in front of the service, and register the API gateway name of the main service as an ALIAS record in Route 53.

**Answer:** B

**NO.18** A company has a VPC that hosts Amazon EC2 instances in a private subnet. The EC2 Instances use a NAT gateway and an internet gateway for internet connectivity to retrieve data from specific internet websites. The company wants to use AWS Network Firewall to filter outbound traffic.

What should a network engineer do to meet these requirements?

- A. 1. Create a firewall in the NAT gateway subnet.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.
3. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.
4. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

**B.** 1. Create a firewall in a new subnet.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.
3. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.
4. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

**C.** 1. Create a firewall in the subnet of the EC2 instances.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.
3. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.
4. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

**D.** 1. Create a firewall in a new subnet.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.
3. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.
4. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

**Answer:** B

**NO.19** A company is deploying a new application on AWS. The application uses dynamic multicasting. The company has five VPCs that are all attached to a transit gateway. Amazon EC2 instances in each VPC need to be able to register dynamically to receive a multicast transmission. How should a network engineer configure the AWS resources to meet these requirements?

**A.** Create a static source multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.

**B.** Create a static source multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.

**C.** Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway.

Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow UDP traffic from the

source to all receivers and to allow UDP traffic that is sent to the multicast group address.

**D.** Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.

**Answer:** C

**NO.20** A company securely connects resources that are in its VPC to a software as a service (SaaS) solution from a SaaS provider. The SaaS solution is hosted in the AWS Cloud and is powered by AWS PrivateLink. The company uses a PrivateLink endpoint to access the SaaS solution behind the SaaS provider's Network Load Balancer (NLB).

The company recently added a new Availability Zone and new subnets to its VPC. A network engineer is unable to deploy a new interface VPC endpoint for the SaaS solution in the new Availability Zone. What is the cause of this problem?

**A.** The CIDR block of the new subnets conflicts with the SaaS provider's CIDR block.

**B.** The enableDnsHostnames attribute and enableDnsSupport attribute were not configured on the new subnets in the new Availability Zone.

**C.** The SaaS provider does not offer the solution in the new Availability Zone and has not configured cross- zone load balancing for the NLB.

**D.** The new subnets are missing a route to the VPC internet gateway.

**Answer:** C

**NO.21** A company's VPC has Amazon EC2 instances that are communicating with AWS services over the public internet. The company needs to change the connectivity so that the communication does not occur over the public internet.

The company deploys AWS PrivateLink endpoints in the VPC. After the deployment of the PrivateLink endpoints, the EC2 instances can no longer communicate at all with the required AWS services.

Which combination of steps should a network engineer take to restore communication with the AWS services?

(Select TWO.)

**A.** In the VPC route table, add a route that has the PrivateLink endpoints as the destination.

**B.** Ensure that the enableDnsSupport attribute is set to True for the VPC. Ensure that each VPC endpoint has DNS support enabled.

**C.** Ensure that the VPC endpoint policy allows communication.

**D.** Create an Amazon Route 53 public hosted zone for all services.

**E.** Create an Amazon Route 53 private hosted zone that includes a custom name for each service.

**Answer:** B C

Explanation:

To use AWS PrivateLink, you need to create interface type VPC endpoints for the services that you want to access privately from your VPC<sup>1</sup>. These endpoints appear as elastic network interfaces (ENIs) with private IPs in your subnets<sup>2</sup>. To enable DNS resolution for these endpoints, you need to set the enableDnsSupport attribute to True for your VPC, and enable DNS support for each endpoint<sup>3</sup>. You also need to ensure that the VPC endpoint policy allows communication between your VPC and the service<sup>4</sup>. You do not need to create any route table entries or Route 53 hosted zones for the endpoints, as they are not required for PrivateLink<sup>5</sup>.

AWS PrivateLink FAQs - Amazon Web Services 2: AWS PrivateLink and service endpoint - Amazon EC2 Overview and Networking Introduction for Telecom Companies 3: VPC Endpoints: Secure and Direct Access to AWS Services 4: AWS PrivateLink and service endpoint - Amazon EC2 Overview and Networking Introduction for Telecom Companies 5: AWS Private Link vs VPC Endpoint - Stack Overflow

**NO.22** A company is running a hybrid cloud environment. The company has multiple AWS accounts as part of an organization in AWS Organizations. The company needs a solution to manage a list of IPv4 on-premises hosts that will be allowed to access resources in AWS. The solution must provide version control for the list of IPv4 addresses and must make the list available to the AWS accounts in the organization.

Which solution will meet these requirements?

**A.** Create a customer-managed prefix list. Add entries for the initial list of on-premises IPv4 hosts. Create a resource share in AWS Resource Access Manager. Add the managed prefix list to the resource share.

Share the resource with the organization.

**B.** Create a customer-managed prefix list. Add entries for the initial list of on-premises IPv4 hosts. Use AWS Firewall Manager to share the managed prefix list with the organization.

**C.** Create a security group. Add inbound rule entries for the initial list of on-premises IPv4 hosts. Create a resource share in AWS Resource Access Manager. Add the security group to the resource share. Share the resource with the organization.

**D.** Create an Amazon DynamoDB table. Add entries for the initial list of on-premises IPv4 hosts. Create an AWS Lambda function that assumes a role in each AWS account in the organization to authorize inbound rules on security groups based on entries from the DynamoDB table.

**Answer:** A

**NO.23** A company is developing a new application that is deployed in multiple VPCs across multiple AWS Regions.

The VPCs are connected through AWS Transit Gateway. The VPCs contain private subnets and public subnets.

All outbound internet traffic in the private subnets must be audited and logged. The company's network engineer plans to use AWS Network Firewall and must ensure that all traffic through Network Firewall is completely logged for auditing and alerting.

How should the network engineer configure Network Firewall logging to meet these requirements?

**A.** Configure Network Firewall logging in Amazon CloudWatch to capture all alerts. Send the logs to a log group in Amazon CloudWatch Logs.

**B.** Configure Network Firewall logging in Network Firewall to capture all alerts and flow logs.

**C.** Configure Network Firewall logging by configuring VPC Flow Logs for the firewall endpoint. Send the logs to a log group in Amazon CloudWatch Logs.

**D.** Configure Network Firewall logging by configuring AWS CloudTrail to capture data events.

**Answer:** B

**NO.24** A company has expanded its network to the AWS Cloud by using a hybrid architecture with multiple AWS accounts. The company has set up a shared AWS account for the connection to its on-premises data centers and the company offices. The workloads consist of private web-based services

for internal use. These services run in different AWS accounts. Office-based employees consume these services by using a DNS name in an on-premises DNS zone that is named `example.internal`. The process to register a new service that runs on AWS requires a manual and complicated change request to the internal DNS. The process involves many teams.

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

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

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

**Answer:** A B D

Explanation:

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

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

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

**NO.25** A company has several production applications across different accounts in the AWS Cloud.

The company operates from the us-east-1 Region only. Only certain partner companies can access the applications. The applications are running on Amazon EC2 instances that are in an Auto Scaling group behind an Application Load Balancer (ALB). The EC2 instances are in private subnets and allow traffic only from the ALB. The ALB is in a public subnet and allows inbound traffic only from partner network IP address ranges over port 80.

When the company adds a new partner, the company must allow the IP address range of the partner network in the security group that is associated with the ALB in each account. A network engineer must implement a solution to centrally manage the partner network IP address ranges.

Which solution will meet these requirements in the MOST operationally efficient manner?

- A.** Create an Amazon DynamoDB table to maintain all IP address ranges and security groups that need to be updated. Update the DynamoDB table with the new IP address range when the company adds a new partner. Invoke an AWS Lambda function to read new IP address ranges and security groups from the DynamoDB table to update the security groups. Deploy this solution in all accounts.
- B.** Create a new prefix list. Add all allowed IP address ranges to the prefix list. Use Amazon EventBridge (Amazon CloudWatch Events) rules to invoke an AWS Lambda function to update security groups whenever a new IP address range is added to the prefix list. Deploy this solution in all accounts.
- C.** Create a new prefix list. Add all allowed IP address ranges to the prefix list. Share the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM). Update security groups to use the prefix list instead of the partner IP address range. Update the prefix list with the new IP address range when the company adds a new partner.
- D.** Create an Amazon S3 bucket to maintain all IP address ranges and security groups that need to be updated. Update the S3 bucket with the new IP address range when the company adds a new partner. Invoke an AWS Lambda function to read new IP address ranges and security groups from the S3 bucket to update the security groups. Deploy this solution in all accounts.

**Answer:** C

Explanation:

Creating a new prefix list and adding all allowed IP address ranges to the prefix list would enable grouping of CIDR blocks that can be referenced in security group rules<sup>3</sup>. Sharing the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM) would enable central management of the partner network IP address ranges<sup>5</sup>. Updating security groups to use the prefix list instead of the partner IP address range would enable simplification of security group rules<sup>3</sup>. Updating the prefix list with the new IP address range when the company adds a new partner would enable automatic propagation of the changes to all security groups that use the prefix list<sup>3</sup>.

**NO.26** Two companies are merging. The companies have a large AWS presence with multiple VPCs and are designing connectivity between their AWS networks. Both companies are using AWS Direct Connect with a Direct Connect gateway. Each company also has a transit gateway and multiple AWS Site-to-Site VPN connections from its transit gateway to on-premises resources. The new solution must optimize network visibility, throughput, logging, and monitoring.

Which solution will meet these requirements?

- A.** Configure a Site-to-Site VPN connection between each company's transit gateway to establish reachability between the respective networks. Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use VPC Reachability Analyzer to monitor connectivity.

- B.** Configure a Site-to-Site VPN connection between each company's transit gateway to establish reachability between the respective networks. Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use AWS Transit Gateway Network Manager to monitor the transit gateways and their respective connections.
- C.** Configure transit gateway peering between each company's transit gateway. Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use VPC Reachability Analyzer to monitor connectivity.
- D.** Configure transit gateway peering between each company's transit gateway. Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use AWS Transit Gateway Network Manager to monitor the transit gateways, their respective connections, and the transit gateway peering link.

**Answer:** D

**NO.27** 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 staging.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:

When a client queries a DNS server for a domain name, the DNS server typically starts by looking for NS records to determine which name servers are authoritative for the domain. The DNS server then queries the authoritative name servers to obtain the information about the domain that the client requested. For example, suppose you own the domain example.com, but you want to delegate control of the subdomain sub.example.com to a different set of name servers. You would create NS records in the example.com zone file that point to the name servers for sub.example.com. This tells DNS servers that the name servers for sub.example.com are authoritative for that subdomain, and they should query those name servers for any requests related to sub.example.com.

com to a different set of name servers. You would create NS records in the example.com zone file that point to the name servers for sub.example.com. This tells DNS servers that the name servers for sub.example.com are authoritative for that subdomain, and they should query those name servers for any requests related to sub.example.com.

**NO.28** A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB).

The ALB is the origin in an Amazon CloudFront distribution. The company wants to implement a

custom authentication system that will provide a token for its authenticated customers.

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

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

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

**Answer:** C

**NO.29** A company recently started using AWS Client VPN to give its remote users the ability to access resources in multiple peered VPCs and resources in the company's on-premises data center. The Client VPN endpoint route table has a single entry of 0.0.0.0/0. The Client VPN endpoint is using a new security group that has no inbound rules and a single outbound rule that allows all traffic to 0.0.0.0/0.

Multiple users report that web search results are showing remote incorrect geographic location information for the users.

Which combination of steps should a network engineer take to resolve this issue with the LEAST amount of service interruption? (Choose three.)

- A.** Switch users to AWS Site-to-Site VPNs.
- B.** Enable the split-tunnel option on the Client VPN endpoint.
- C.** Add routes for the peered VPCs and for the on-premises data center to the Client VPN route table.
- D.** Remove the 0.0.0.0/0 outbound rule from the security group that the Client VPN endpoint uses.
- E.** Delete and recreate the Client VPN endpoint in a different VPC.
- F.** Remove the 0.0.0.0/0 entry from the Client VPN endpoint route table.

**Answer:** B C F