[2025-November-NewBraindump2go ANS-C01 PDF Free Updated[Q143-Q176

2025/November Latest Braindump2go ANS-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go ANS-C01 Real Exam Questions!QUESTION 143A company has set up hybrid connectivity between its VPCs and its on-premises data center. The company has the on-premises.example.com subdomain configured at its DNS server in the on-premises data center. The company is using the aws.example.com subdomain for workloads that run on AWS across different VPCs and accounts. Resources in both environments can access each other by using IP addresses. The company wants workloads in the VPCs to be able to access resources on premises by using the on-premises.example.com DNS names.Which solution will meet these requirements with MINIMUM management of resources?A. Create an Amazon Route 53 Resolver outbound endpoint. Configure a Resolver rule that conditionally forwards DNS queries for on-premises.example.com to the on-premises DNS server. Associate the rule with the VPCs.B. Create an Amazon Route 53 Resolver inbound endpoint and a Resolver outbound endpoint. Configure a Resolver rule that conditionally forwards DNS queries for on-premises.example.com to the on-premises DNS server. Associate the rule with the VPCs.C. Launch an Amazon EC2 instance. Install and configure BIND software to conditionally forward DNS queries for on-premises.example.com to the on-premises DNS server. Configure the EC2 instance's IP address as a custom DNS server in each VPC.D. Launch an Amazon EC2 instance in each VPC. Install and configure BIND software to conditionally forward DNS queries for on-premises.example.com to the on-premises DNS server. Configure the EC2 instance's IP address as a custom DNS server in each VPC.Answer: AExplanation:We need an outbound endpoint because we want to resolve it with an on-premises DNS query.QUESTION 144A 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: BExplanation:Acceleration is only supported for Site-to-Site VPN connections that are attached to a transit gateway. Virtual private gateways do not support accelerated VPN connections.
https://docs.aws.amazon.com/vpn/latest/s2svpn/accelerated-vpn.htmlQUESTION 145A company is moving its record-keeping application to the AWS Cloud. All traffic between the company's on-premises data center and AWS must be encrypted at all times and at every transit device during the migration.The application will reside across multiple Availability Zones in a single AWS Region. The application will use existing 10 Gbps AWS Direct Connect dedicated connections with a MACsec capable port. A network engineer must ensure that the Direct Connect connection is secured accordingly at every transit device.The network engineer creates a Connection Key Name and Connectivity Association Key (CKN/CAK) pair for the MACsec secret key.Which combination of additional steps should the network engineer take to meet the requirements? (Choose two.)A. Configure the on-premises router with the MACsec secret key.B. Update the connection's MACsec encryption mode to must_encrypt. Then associate the CKN/CAK pair with the connection.C. Update the connection's MACsec encryption mode to should encrypt. Then associate the CKN/CAK pair with the connection.D. Associate the CKN/CAK pair with the connection. Then update the connection's MACsec encryption mode to must_encrypt.E. Associate the CKN/CAK pair with the connection. Then update the connection's MACsec encryption mode to should_encrypt.Answer: ADExplanation:According to AWS, you need to do the following 4 steps in order.1. Create a new connection with MACsec support2. Associate the CKN/CAK with the connection3. Verify the connection status4. Migrate traffic to new connection as appropriateWhen you first create the DX connection, the default encryption mode is should encrypt. You need to update it to must encrypt in step 3. There's no way to specify that during the creation of DX.
https://aws.amazon.com/blogs/networking-and-content-delivery/adding-macsec-security-to-aws-direct-connect-connections/
QUESTION 146A 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: DExplanation:The VPC route to send to on-premises is sent by entering the allowed prefix value of DXGW. Since only one remaining frame is used for route information, it is necessary to aggregate two routes.QUESTION 147Two 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: DExplanation:Transit gateway peering will allow the communication between all networks. To monitor the overall infrastructure, AWS Transit Gateway Network Manager is utilized for this purpose.https://aws.amazon.com/transit-gateway/network-manager/QUESTION 148A company has a single VPC in the us-east-1 Region. The company is planning to set up a new VPC in the us-east-2 Region. The existing VPC has an AWS Site-to-Site VPN connection to the company's on-premises environment and uses a virtual private gateway.A network engineer needs to implement a solution to establish connectivity between the existing VPC and the new VPC. The solution also must implement support for IPv6 for the new VPC. The company has new on-premises resources that need to connect to VPC resources by using IPv6 addresses.Which solution will meet these requirements?A. Create a new virtual private gateway in us-east-1. Attach the new virtual private gateway to the new VPC. Create two new Site-to-Site VPN connections to the new virtual private gateway with IPv4 and IPv6 support. Configure routing between the VPCs by using VPC peering.B. Create a transit gateway in us-east-1 and in us-east-2. Attach the existing VPC and the new VPC to each transit gateway. Create a new Site-to-Site VPN connection to each transit gateway with IPv4 and IPv6 support. Configure transit gateway peering. Configure routing between the VPCs and the on-premises environment.C. Create a new virtual private gateway in us-east-2. Attach the new virtual private gateway to the new VPCCreate two new Site-to-Site VPN connections to the new virtual private gateway with IPv4 and IPv6 support. Configure routing between the VPCs by using VPC peering.D. Create a transit gateway in us-east-1. Attach the existing VPC and the new VPC to the transit gateway. Create two new Site-to-Site VPN connections to the transit gateway with IPv4 and IPv6 support. Configure transit gateway peering. Configure routing between the VPCs and the on-premises environment.Answer: BExplanation:Transit gateway attachment can only be in the same region as the TGW itself.QUESTION 149A 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 "*" 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: DExplanation:If the domain name in a query doesn't match the domain name in any other rules, Resolver forwards the query based on the settings in the autodefined "." (dot) rule. The dot rule applies to all domain names except some AWS internal domain names and record names in private hosted zones. https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-overview-DSN-queries-to-vpc.html#resolver-overview-forward-vpc-to-network-domain-name-matchesQUESTION 150A 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: ACExplanation:A transit gateway supports an MTU of 8500 bytes for traffic between VPCs, AWS Direct Connect, Transit Gateway Connect, and peering attachments. Traffic over VPN connections can have an MTU of 1500 bytes.QUESTION 151An 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: BExplanation:CloudFront is designed to handle HTTP protocol meanwhile Global Accelerator is best used for both HTTP and non-HTTP protocols such as TCP and UDP. https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.htmlQUESTION 152 A company is deploying a new stateless web application on AWS. The web application will run on Amazon EC2 instances in private subnets behind an Application Load Balancer. The EC2 instances are in an Auto Scaling group. The web application has a stateful management application for administration that will run on EC2 instances that are in a separate Auto Scaling group.The company wants to access the management application by using the same URL as the web application, with a path prefix of/management. The protocol, hostname, and port number must be the same for the web application and the management application. Access to the management application must be restricted to the company's on-premises IP address space. An SSL/TLS certificate from AWS Certificate Manager (ACM) will protect the web application.Which combination of steps should a network engineer take to meet these requirements? (Choose two.)A.   Insert a rule for the load balancer HTTPS listener. Configure the rule to check the path-pattern condition type for the /management prefix and to check the source-ip condition type for the on-premises IP address space. Forward requests to the management application target group if there is a match. Edit the management application target group and enable stickiness.B.   Modify the default rule for the load balancer HTTPS listener. Configure the rule to check the path-pattern condition type for the /management prefix and to check the source-ip condition type for the on-premises IP address space. Forward requests to the management application target group if there is not a match. Enable group-level stickiness in the rule attributes.C.   Insert a rule for the load balancer HTTPS listener. Configure the rule to check the path-pattern condition type for the /management prefix and to check the X-Forwarded-For HTTP header for the on-premises IP address space. Forward requests to the management application target group if there is a match. Enable group-level stickiness in the rule attributes.D.   Modify the default rule for the load balancer HTTPS listener. Configure the rule to check the path-pattern condition type for the /management prefix and to check the source-ip condition type for the on-premises IP address space. Forward requests to the web application target group if there is not a match.E. Forward all requests to the web application target group. Edit the web application target group and disable stickiness.Answer: ADExplanation:Default rule does not need stickiness because it is stateless.QUESTION 153A company deploys a software solution on Amazon EC2 instances that are in a cluster placement 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: BCExplanation: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/network_mtu.htmlQUESTION 154A company has users who work from home. The company wants to move these users to Amazon WorkSpaces for additional security visibility.The company has deployed WorkSpaces in its own AWS account in VPC A. A network engineer decides to provide the security visibility by using two firewall appliances behind a Gateway Load Balancer (GWLB). The network engineer provisions another VPC, VPC B, in a separate account and deploys the two firewall appliances in separate Availability Zones.What should the network engineer do to configure the network connectivity for this solution?A. Create a GWLB in VPC A with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the WorkSpaces account to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the default route to the VPC endpoint.B. Create a GWLB in VPC B with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the WorkSpaces account to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the default route to the GWLB endpoint.C. Create a GWLB in VPC B with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the WorkSpaces account to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the WorkSpaces subnet to the VPC endpoint.D. Create a GWLB in VPC B with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the account that contains the firewall appliances to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the default route to the VPC endpoint.Answer: BExplanation:Using AWS PrivateLink, GWLB Endpoint routes traffic to GWLB. Traffic is routed securely over Amazon network without any additional configuration.QUESTION 155A 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: CExplanation:Return traffic was blocked by NACL, outbound should be allowed.QUESTION 156A company needs to temporarily scale out capacity for an on-premises application and wants to deploy new servers on Amazon EC2 instances. A network engineer must design the networking solution for the connectivity and for the application on AWS.The EC2 instances need to share data with the existing servers in the on-premises data center. The servers must not be accessible from the internet. All traffic

to the internet must route through the firewall in the on-premises data center. The servers must be able to access a third-party web application.Which configuration will meet these requirements?A. Create a VPC that has public subnets and private subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a NAT gateway in a public subnet. Create a route table, and associate the public subnets with the route table. Add a default route to the internet gateway. Create a route table, and associate the private subnets with the route table. Add a default route to the NAT gateway. Add routes for the data center subnets to the virtual private gateway. Deploy the application to the private subnets.B. Create a VPC that has private subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a route table, and associate the private subnets with the route table. Add a default route to the virtual private gateway. Deploy the application to the private subnets.C. Create a VPC that has public subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a route table, and associate the public subnets with the route table. Add a default route to the internet gateway. Add routes for the on-premises data center subnets to the virtual private gateway. Deploy the application to the public subnets.D. Create a VPC that has public subnets and private subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a route table, and associate the public subnets with the route table. Add a default route to the internet gateway. Create a route table, and associate the private subnets with the route table. Add routes for the on-premises data center subnets to the virtual private gateway. Deploy the application to the private subnets.Answer: BExplanation:You don't need to a private subnet as you should only be able to get to the instances from on prem, also you don't need a public subnet with a nat gateway as internet traffic goes through on prem firewall.QUESTION 157A company is deploying a web application into two AWS Regions. The company has one VPC in each Region. Each VPC has three Amazon EC2 instances as web servers behind an Application Load Balancer (ALB). The company already has configured an Amazon Route 53 public hosted zone for example.com. Users will access the application by using the fully qualified domain name (FQDN) of app.example.com.The company needs a DNS solution that allows global users to access the application. The solution must route the users' requests to the Region that provides the lowest response time. The solution must fail over to the Region that provides the next-lowest response time if the application is unavailable in the initially intended Region.Which solution will meet these requirements?A. For each ALB, create an A record that has a geolocation routing policy to route app.example.com to the IP addresses of the ALB. Configure a Route 53 HTTP health check that monitors each ALB by IP address. Associate the health check with the A records.B. Create an A record that has a geolocation routing policy to route app.example.com to the IP addresses for both ALBs. Configure a Route 53 health check that monitors TCP port 80 for each ALB by IP address. Associate the health check with the A records.C. Create an A record that has a latency-based routing policy to route app.example.com as an alias to one of the ALBs. Configure a Route 53 health check that monitors TCP port 80 for each ALB by IP address. Associate the health check with the A records.D. For each ALB, create an A record that has a latency-based routing policy to route app.example.com as an alias to the ALB. Set the value for Evaluate Target Health to Yes for the records.Answer: DQUESTION 158A consulting company manages AWS accounts for its customers. One of the company's customers needs to add intrusion prevention for its environment without having to re-architect the environment. The customer's environment includes five VPCs in two AWS Regions in the United States. VPC-to-VPC connectivity is achieved through VPC peering. The customer does not plan to increase the number of VPCs within the next 2 years. The solution must accommodate unencrypted traffic.Which solution will meet these requirements?A. Configure VPC security groups and network ACLs.B. Use an AWS Network Firewall centralized deployment model in each VPC.C. Use an AWS Network Firewall distributed deployment model in each VPC.D. Deploy AWS Shield in each VPC.Answer: CExplanation:You can use the same model for inspection of traffic to other AWS Regions using AWS Transit Gateway Inter-Region Peering feature as shown in Figure 8. Remote AWS Regions are treated as spokes.
https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/QUESTION 159A 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: DExplanation:Per VPN connection, you can achieve 1.25 Gbps of throughput and 140,000 packets per second. When terminating the VPN connections in the Transit Gateway, you can use Equal Cost Multi-Path (ECMP) routing to get a higher VPN bandwidth by aggregating multiple VPN tunnels. To use ECMP, you need to configure dynamic routing in the VPN connections ? ECMP is not supported using static routing. https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-transit-gateway-vpn.htmlQUESTION 160A company has business operations in the United States and in Europe. The company's public applications are running on AWS and use three transit gateways. The transit gateways are located in the us-west-2, us-east-1, and eu-central-1 Regions. All the transit gateways are connected to each other in a full mesh configuration.The company accidentally removes the route to the eu-central-1 VPCs from the us-west-2 transit gateway route table. The company also accidentally removes the route to the us-west-2 VPCs from the eu-central-1 transit gateway route table.How can a network engineer identify the misconfiguration with the LEAST operational overhead?A. Use the Route Analyzer feature for AWS Transit Gateway Network Manager.B. Use the AWSSupport-SetupIPMonitoringFromVPC AWS Systems Manager Automation runbook. Push network telemetry data to Amazon CloudWatch Logs for analysis.C. Use VPC flow logs in eu-central-1 and us-west-2 to analyze the missing routes.D. Use Amazon VPC Traffic Mirroring in eu-central-1 or us-west-2 to take packet captures and troubleshoot the connectivity issues.Answer: AExplanation:https://docs.aws.amazon.com/network-manager/latest/tgwnm/route-analyzer.htmlQUESTION 161A marketing company is using hybrid infrastructure through AWS Direct Connect links and a software-defined wide area network (SD-WAN) overlay to connect its branch offices. The company connects multiple VPCs to a third-party SD-WAN appliance transit VPC within the same account by using AWS Site-to-Site VPNs.The company is planning to connect more VPCs to the SD-WAN appliance transit VPC. However, the company faces challenges of scalability, route table limitations, and higher costs with the existing architecture. A network engineer must design a solution to resolve these issues and remove dependencies.Which solution will meet these requirements with the LEAST amount of operational overhead?A. Configure a transit gateway to attach the VPCs. Configure a Site-to-Site VPN connection between the transit gateway and the third-party SD-WAN appliance transit VPC. Use the SD-WAN overlay links to connect to the branch offices.B. Configure a transit gateway to attach the VPCs. Configure a transit gateway Connect attachment for the third-party SD-WAN appliance transit VPC. Use transit gateway Connect native integration of SD-WAN virtual hubs with AWS Transit Gateway.C. Configure a transit gateway to attach the VPCs. Configure VPC peering between the VPCs and the third-party SD-WAN appliance transit VPUse the SD-WAN overlay links to connect to the branch offices.D. Configure VPC peering between the VPCs and the third-party SD-WAN appliance transit VPC. Use transit gateway Connect native integration of SD-WAN virtual hubs with AWS Transit Gateway.Answer: BExplanation: https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-transit-gateway-sd-wan.htmlQUESTION 162A 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 Explanation:https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.htmlQUESTION 163A company's application is deployed on Amazon EC2 instances in a single VPC in an AWS Region. The EC2 instances are running in two Availability Zones. The company decides to use a fleet of traffic inspection instances from AWS Marketplace to inspect traffic between the VPC and the internet. The company is performing tests before the company deploys the architecture into production. The fleet is located in a shared inspection VPC behind a Gateway Load Balancer (GWLB). To minimize the cost of the solution, the company deployed only one inspection instance in each Availability Zone that the application uses.During tests, a network engineer notices that traffic inspection works as expected when the network is stable. However, during maintenance of the inspection instances, the internet sessions time out for some application instances. The application instances are not able to establish new sessions.Which combination of steps will remediate these issues? (Choose two.)A. Deploy one inspection instance in the Availability Zones that do not have inspection instances deployed.B. Deploy one additional inspection instance in each Availability Zone where the inspection instances are deployed.C. Enable the cross-zone load balancing attribute for the GWLB.D. Deploy inspection instances in an Auto Scaling group. Define a scaling policy that is based on CPU load.E. Attach the GWLB to all Availability Zones in the Region.Answer: BCExplanation: https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-inspection-architecture-with-aws-gateway-load-balancer-and-aws-transit-gateway/QUESTION 164A company has developed a new web application on AWS. The application runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate behind an Application Load Balancer (ALB) in the us-east-1 Region. The application uses Amazon Route 53 to host the DNS records for the domain. The content that is served from the website is mostly static images and files that are not updated frequently. Most of the traffic to the website from end users will originate from the United States. Some traffic will originate from Canada and Europe.A network engineer needs to design a solution that will reduce latency for end users at the lowest cost. The solution also must ensure that all traffic is encrypted in transit until the traffic reaches the ALB.Which solution will meet these requirements?A. Configure the ALB to use an AWS Global Accelerator accelerator in us-east-1. Create a secure HTTPS listener. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to route to the DNS name that is assigned to the accelerator for the ALB.B. Configure the ALB to use a secure HTTPS listener. Create an Amazon CloudFront distribution. Set the origin domain name to point to the DNS record that is assigned to the ALConfigure the CloudFront distribution to use an SSL certificate. Set all behaviors to force HTTPS. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to route to the DNS name that is assigned to the ALB.C. Configure the ALB to use a secure HTTPS listener. Create an Amazon CloudFront distribution. Set the origin domain name to point to the DNS record that is assigned to the ALB. Configure the CloudFront distribution to use an SSL certificate and redirect HTTP to HTTPS. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to route to the CloudFront distribution.D. Configure the ALB to use an AWS Global Accelerator accelerator in us-east-1. Create a secure HTTPS listener. Create a second application stack on Amazon ECS on Fargate in the eu-west-1 Region. Create another secure HTTPS listener. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to use a latency-based routing policy to route to the DNS name that is assigned to the accelerator for the ALBs.Answer: CExplanation:Route 53 record points to Cloudfront default DNS name.QUESTION 165A company deploys an internal website behind an Application Load Balancer (ALB) in a VPC. The VPC has a CIDR block of 172.31.0.0/16. The company creates a private hosted zone for the domain example.com for the website in Amazon Route 53. The company establishes an AWS Site-to-Site VPN connection between its office network and the VPC.A network engineer needs to set up a DNS solution so that employees can visit the internal webpage by accessing a private domain URL (https://example.com) from the office network.Which combination of steps will meet this requirement? (Choose two.)A. Create an alias record that points to the ALB in the Route 53 private hosted zone.B. Create a CNAME record that points to the ALB internal domain in the Route 53 private hosted zone.C. Create a Route 53 Resolver inbound endpoint. On the office DNS server, configure a conditional forwarder to forward the DNS queries to the Route 53 Resolver inbound endpoint.D. Create a Route 53 Resolver outbound endpoint. On the office DNS server, configure a conditional forwarder to forward the DNS queries to the Route 53 Resolver outbound endpoint.E. On the office DNS server, configure a conditional forwarder for the private domain to the VPC DNS at 172.31.0.2.Answer: ACExplanation:You can't create a CNAME record that has the same name as the hosted zone (the zone apex). This is true both for hosted zones for domain names (example.com) and for hosted zones for subdomains (zenith.example.com). https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.htmlQUESTION 166A company is deploying AWS Cloud WAN with edge locations in the us-east-1 Region and the ap-southeast-2 Region. Individual AWS Cloud WAN segments are configured for the development environment, the production environment, and the shared services environment at each edge location. Many new VPCs will be deployed for the environments and will be configured as attachments to the AWS Cloud WAN core network.The company's network team wants to ensure that VPC attachments are configured for the correct segment. The network team will tag the VPC attachments by using the Environment key with a value of the corresponding environment segment name. The segment for the production environment in us-east-1 must require acceptance for attachment requests. All other attachment requests must not require acceptance.Which solution will meet these requirements?A. Create a rule with a number of 100 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the "or" value. Include conditions that require a tag:Environment value of Production or a Region value of us-east-1. Create a rule with a number of 200 that does not require acceptance to map any tag:Environment values to their respective segments.B. Create a rule with a number of 100 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the

"and" value. Include conditions that require a tag:Environment value of Production and a Region value of us-east-1. Create a rule with a number of 200 that does not require acceptance to map any tag.Environment values to their respective segments.C. Create a rule with a number of 100 that does not require acceptance to map any tag:Environment values to their respective segments. Create a rule with a number of 200 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the "and" value. Include conditions that require a tag:Environment value of Production and a Region value of us-east-1.D. Create a rule with a number of 100 that does not require acceptance to map any tag:Environment values to their respective segments. Create a rule with a number of 200 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the "or" value. Include conditions that require a tag:Environment value of Production or a Region value of us-east-1.Answer: B Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/achieving-traffic-segmentation-in-multi-aws-region-environments-using-aws-transit-gateway-and-aws-cloud-wan/QUESTION 167A company is migrating applications from a data center to AWS. Many of the applications will need to exchange data with the company's on-premises mainframe.The company needs to achieve 4 Gbps transfer speeds to meet peak traffic demands. A network engineer must design a highly available solution that maximizes resiliency. The solution must be able to withstand the loss of circuits or routers.Which solution will meet these requirements?A. Order four 10 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate one connection from each Direct Connect location to a router at the company location. Terminate the other connection from each Direct Connect location to a different router at the company location.B. Order two 10 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate the connection from each Direct Connect location to a different router at the company location.C. Order four 1 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate one connection from each Direct Connect location to a router at the company location. Terminate the other connection from each Direct Connect location to a different router at the company location.D. Order two 1 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate the connection from each Direct Connect location to a different router at the company location.Answer: BQUESTION 168A company has 10 web server Amazon EC2 instances that run in an Auto Scaling group in a production VPC. The company has 10 other web servers that run in an on-premises data center. The company has a 10 Gbps AWS Direct Connect connection between the on-premises data center and the production VPC.The company needs to implement a load balancing solution that receives HTTPS traffic from thousands of external users. The solution must distribute the traffic across the web servers on AWS and the web servers in the on-premises data center. Regardless of the location of the web servers, HTTPS requests must go to the same web server throughout the entire session.Which solution will meet these requirements?A. Create a Network Load Balancer (NLB) in the production VPC. Create a target group. Specify ip as the target type. Register the EC2 instances and the on-premises servers with the target group Enable connection draining on the NLBB. Create an Application Load Balancer (ALB) in the production VPC. Create a target group Specify ip as the target type. Register the EC2 instances and the on-premises servers with the target group. Enable application-based session affinity (sticky sessions) on the ALB.C. Create a Network Load Balancer (NLB) in the production VPCreate a target group. Specify instance as the target type. Register the EC2 instances and the on-premises servers with the target group. Enable session affinity (sticky sessions) on the NLB.D. Create an Application Load Balancer (ALB) in the production VPC. Create a target group. Specify instance as the target type Register the EC2 instances and the on-premises servers with the target group Enable application-based session affinity (sticky sessions) on the ALB.Answer: BExplanation:ALB support on prem's ip address as a target group, and you need session affinity for this.
https://aws.amazon.com/blogs/aws/new-application-load-balancing-via-ip-address-to-aws-on-premises-resources/QUESTION 169A 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 AWS side 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 without specifying the IP address of the customer gateway device.Answer: BCFExplanation:An IP address is not required when you are using a private certificate from AWS Private Certificate Authority.https://docs.aws.amazon.com/vpn/latest/s2svpn/cgw-options.htmlQUESTION 170A company's

AWS environment has two VPCs. VPC A has a CIDR block of 192.168.0.0/16. VPC B has a CIDR block of 10.0.0.0/16. Each VPC is deployed in a separate AWS Region. The company has remote users who work outside the company's offices. These users need to connect to an application that is running in the VPCs.Traffic to and from the VPCs over the internet must be encrypted. A network engineer must set up connectivity between the remote users and the VPCs.Which combination of steps should the network engineer take to meet these requirements with the LEAST management overhead? (Choose three.)A. Establish an AWS Site-to-Site VPN connection between VPC A and VPC B.B. Establish a VPC peering connection between VPC A and VPC B.C. Create an AWS Client VPN endpoint in VPC A and VPC B Add an authorization rule to grant access to VPC A and VPC B.D. Create an AWS Client VPN endpoint in VPC A Add an authorization rule to grant access to VPC A and VPC B.E. Add a route to the AWS Client VPN endpoint's route table to direct traffic to VPC B.F. Add a route to the AWS Client VPN endpoint's route table to direct traffic to VPC A.Answer: BDEExplanation:The procedure for allowing access to a peered VPC outlined below, is only required if the Client VPN endpoint was configured for split-tunnel mode. In full-tunnel mode, access to the peered VPC is allowed by default. https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/scenario-peered.htmlQUESTION 171A company uses Amazon Route 53 to register a public domain, example.com, in an AWS account. A central services group manages the account. The company wants to create a subdomain, test.example.com, in another AWS account to offer name services for Amazon EC2 instances that are hosted in the account. The company does not want to migrate the parent domain to the subdomain account.A network engineer creates a new Route 53 hosted zone for the subdomain in the second account.Which combination of steps must the network engineer take to complete the task? (Choose two.)A. Add records for the hosts of the new subdomain to the new Route 53 hosted zone.B. Update the DNS service for the parent domain by adding name server (NS) records for the subdomain.C. Update the DNS service for the subdomain by adding name server (NS) records for the parent domain.D. Create an alias record from the parent domain that points to the hosted zone for the subdomain in the second account.E. Add a start of authority (SOA) record in the parent domain for the subdomain.Answer: ABExplanation:https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/CreatingNewSubdomain.html QUESTION 172An IoT company collects data from thousands of sensors that are deployed in the Unites States and South Asia. The sensors use a proprietary communication protocol that is built on UDP to send the data to a fleet of Amazon EC2 instances. The instances are in an Auto Scaling group and run behind a Network Load Balancer (NLB). The instances, Auto Scaling group, and NLB are deployed in the us-west-2 Region.Occasionally, the data from the sensors in South Asia gets lost in transit over the internet and does not reach the EC2 instances.Which solutions will resolve this issue? (Choose two.)A. Use AWS Global Accelerator with the existing NLB.B. Create an Amazon CloudFront distribution. Specify the existing NLB as the origin.C. Create a second deployment of the EC2 instances and the NLB in the ap-south-1 Region. Use an Amazon Route 53 latency routing policy to resolve to the Region that provides the least latency.D. Create a second deployment of the EC2 instances and the NLB in the ap-south-1 Region. Use an Amazon Route 53 failover routing policy to resolve to an alternate Region in case packets are dropped.E. Turn on enhanced networking on the EC2 instances by using the most recent Elastic Network Adapter (ENA) drivers.Answer: ACQUESTION 173A 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 Kinesis 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 Kinesis Data 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 Kinesis Data Firehose to load the logs into the appliance.Answer: BExplanation:You can use the following resources as traffic mirror targets:- Network interfaces of type interface- Network Load Balancers- Gateway Load Balancer endpoints https://docs.aws.amazon.com/vpc/latest/mirroring/traffic-mirroring-targets.htmlQUESTION 174A company has two AWS Direct Connect links. One Direct Connect link terminates in the us-east-1 Region, and the other Direct Connect link terminates in the af-south-1 Region. The company is using BGP to exchange routes with AWS.How should a network engineer configure BGP to ensure that af-south-1 is used as a secondary link to AWS?A. On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7100On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7300On the Direct Connect BGP peer to us-east-1, set the local preference value to 200On the Direct Connect BGP peer to af-south-1, set the local preference value to 50B. On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7300On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7100On the Direct Connect BGP peer to us-east-1, set the local preference value to 200On the Direct Connect BGP peer to af-south-1, set the local preference value to 50C. On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7100On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7300On the Direct Connect BGP peer to us-east-1, set the local preference value to 50On the Direct Connect BGP peer to af-south-1, set the local preference value to 200D. On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7300On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7100On the Direct Connect BGP peer to us-east-1, set the local preference value to 50On the Direct Connect BGP peer to af-south-1, set the local preference value to 200Answer: BExplanation: https://docs.aws.amazon.com/directconnect/latest/UserGuide/routing-and-bgp.htmlQUESTION 175A team of infrastructure engineers wants to automate the deployment of Application Load Balancer (ALB) components by using the AWS Cloud Development Kit (AWS CDK). The CDK application must deploy an infrastructure stack that is reusable and consistent across multiple environments, AWS Regions, and AWS accounts.The lead network architect on the project has already bootstrapped the target accounts. The lead network architect also has deployed core network components such as VPCs and Amazon Route 53 private hosted zones across the multiple environments and Regions. The infrastructure engineers must design the ALB components in the CDK application to use the existing core network components.Which combination of steps will meet this requirement with the LEAST manual effort between environment deployments? (Choose two.)A. Design the CDK application to read AWS CloudFormation parameters for the values that vary across environments and Regions. Reference these variables in the CDK stack for resources that require the variables.B. Design the CDK application to read environment variables that contain account and Region details at runtime. Use these variables as properties of the CDK stack. Use context methods in the CDK stack to retrieve variable values.C. Create a dedicated account for shared application services in the multi-account environment. Deploy a CDK pipeline to the dedicated account. Create stages in the pipeline that deploy the CDK application across different environments and Regions.D. Write a script that automates the deployment of the CDK application across multiple environments and Regions. Distribute the script to engineers who are working on the project.E. Use the CDK toolkit locally to deploy stacks to each environment and Region. Use the --context flag to pass in variables that the CDK application can reference at runtime.Answer: BC Explanation:https://docs.aws.amazon.com/cdk/v2/guide/environments.htmlyou can use environment variables to pass in values that vary across environments and Regions. You can use the --context flag when running cdk deploy to set environment variables for the CDK application. You can also use the CDK_DEFAULT_ACCOUNT and CDK_DEFAULT_REGION environment variables provided by the AWS CDK CLI to specify the target account and Region for deployment.QUESTION 176A company has critical VPC workloads that connect to an on-premises data center through two redundant active-passive AWS Direct Connect connections. However, a recent outage on one Direct Connect connection revealed that it takes more than a minute for traffic to fail over to the secondary Direct Connect connection. The company wants to reduce the failover time from minutes to seconds.Which solution will provide the LARGEST reduction in the BGP failover time?A. Reduce the BGP hold-down timer that is configured on the BGP sessions on the Direct Connect connection VIFs.B. Configure an Amazon CloudWatch alarm for the Direct Connect connection state to invoke an AWS Lambda function to fail over the traffic.C. Configure Bidirectional Forwarding Detection (BFD) on the Direct Connect connections on the AWS side.D. Configure Bidirectional Forwarding Detection (BFD) on the Direct Connect connections on the on-premises router.Answer: DExplanation: https://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.htmlBy enabling BFD on both sides of the Direct Connect connection, you can reduce the BGP failover time from minutes to seconds. BFD allows the BGP neighbor relationship to be quickly torn down when a failure is detected on the Direct Connect connection. Otherwise, by default, BGP waits for three keep-alives to fail at a hold-down time of 90 seconds.Resources From:1.2025 Latest Braindump2go ANS-C01 Exam Dumps (PDF & VCE) Free Share:**https://www.braindump2go.com/ans-c01.html**2.2025 Latest Braindump2go ANS-C01 PDF and ANS-C01 VCE Dumps Free Share:**https://drive.google.com/drive/folders/1l_8zUaGGHOED0OZGVvaaxcDC_V1R6CWN?usp=sharing3.2025 Free Braindump2go ANS-C01 Exam Questions Download: https://www.braindump2go.com/free-online-pdf/ANS-C01-VCE-Dumps(143-176).pdfFree Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!**