Free Amazon ANS-C01 Practice Questions

Shared by Carlson on 12-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company is establishing hybrid cloud connectivity from an on-premises environment to AWS in the us-east-1 Region. The company is using a 10 Gbps AWS Direct Connect dedicated connection. The company has two accounts in AWS. Account A has transit gateways in four AWS Regions. Account has transit gateways in three Regions. The company does not plan to expand.

To meet security requirements the company's accounts must have separate cloud infrastructure.

Which solution will meet these requirements MOST cost-effectively?

Options:

A- Create one Direct Connect gateway in us-east-1. Use AWS Resource Access Manager (AWS RAM) to share the Direct Connect gateway with each account. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway. Create a transit VIF for Account B. Associate the three transit gateways in Account to the Direct Connect gateway.
B- Create one Direct Connect gateway in us-east-1 for Account A. Create a second Direct Connect gateway in us-east-1 for Account B. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway in Account A. Create a transit VIF for Account B. Associate the three transit gateways in Account to the Direct Connect gateway in Account .
C- Create one Direct Connect gateway in us-east-1. Use AWS Resource Access Manager (AWS RAM) to share the Direct Connect gateway with each account. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway. Order a new 10 Gbps Direct Connect dedicated connection for Account B. Create a transit VIF on the new Direct Connect connection for Account B. Associate the three transit gateways in Account to the Direct Connect gateway.
D- Create one Direct Connect gateway in us-east-1 for Account A. Create a second Direct Connect gateway in us-east-1 for Account B. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway in Account A. Order a new 10 Gbps Direct Connect dedicated connection for Account . Create a transit VIF on the new Direct Connect connection for Account . Associate the three transit gateways in Account to the Direct Connect gateway in Account .

Answer:

A

Explanation:

The most cost-effective and scalable solution is to create a single Direct Connect gateway in us-east-1, and use AWS Resource Access Manager (AWS RAM) to share the Direct Connect

gateway between Account A and Account B. This approach avoids the need for multiple Direct Connect connections and allows both accounts to share the same connection, which is a more cost-efficient solution compared to creating separate connections for each account.

Transit VIFs (Virtual Interfaces) will be created for both Account A and Account B, and each account's respective transit gateways will be associated with the same Direct Connect gateway. This solution allows both accounts to access AWS resources in the most efficient manner.

Question 2

Question Type: MultipleChoice

A bank built a new version of its banking application in AWS using containers that content to an on-premises database over VPN connection. This application version requires users to also update their client application. The bank plans to deprecate the earlier client version. However, the company wants to keep supporting earlier clients through their on-premises version of the application to serve a small portion of the customers who haven't yet upgraded.

What design will allow the company to serve both newer and earlier clients in the MOST efficient way?

Options:

A- Use an Amazon Route 53 multivalue answer routing policy to route older client traffic to the on-premises application version and the rest of the traffic to the new AWS based version.
B- Use a Classic Load Balancer for the new application. Route all traffic to the new application by using an Elastic Load Balancing (ELB) load balancer DNS. Define a user-agent-based rule on the backend servers to redirect earlier clients to the on-premises application.
C- Use an Application Load Balancer for the new application. Register both the new and earlier applications as separate target groups and use path-based routing to route traffic based on the application version.
D- Use an Application Load Balancer for the new application. Register both the new and earlier application backends as separate target groups. Use header-based routing to route traffic based on the application version.

Answer:

D

Question 3

Question Type: MultipleChoice

A company is planning to migrate to AWS and use multiple VPCs in multiple AWS Regions. A network engineer must connect the eu-west-1 and eu-central-1 Regions to the company headquarters and branch office, respectively.

The network engineer created a production VPC, named Prod A, with a CIDR block of 10.0.0.0/16. Prod A runs in an account in eu-west-1. The network engineer then created another production VPC, named Prod B, with a CIDR block of 10.1.0.0/16. Prod runs in a different account in eu-central-1.

The network engineer performed the following steps to try to achieve the required connectivity:

1. Created one transit gateway in each Region

2. Shared and accepted the transit gateways with the production accounts in both Regions

3. Configured the peering attachment between both transit gateways

4. Attached both VPCs to the respective Region transit gateway

5. Created both transit gateway route tables and associated the attachments with the route tables

6. Configured a static route in both transit gateway route tables to send traffic to the remote VPC in the other Region

7. Activated route propagation on the VPC route tables in each Region

After the configuration, the network engineer tried to connect from Prod A to Prod B. However, the connection was unsuccessful.

What should the network engineer do to achieve the required connectivity?

Options:

A- Modify the IP address of the peering attachment to a wider range.
B- Delete the static routes that were in the transit gateway route table to send traffic to the remote VPC and enable route propagation instead.
C- Create a new route destined to 10.0.0.0/8 in both production VPC route tables with the Region transit gateway as the target.
D- Modify the transit gateway route tables from the production accounts to propagate routes dynamically between the production VPCs.

Answer:

C

Question 4

Question Type: MultipleChoice

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

Which solution will meet these requirements MOST cost-effectively?

Options:

A- Create a new VPC in the Local Zone. Attach all the VPCs to a transit gateway. Configure routing for the transit gateway and the VPCs. Deploy instances in the new VPC.
B- Deploy a third-party appliance in a new VPC in the Region. Create a new VPC in the Local Zone. Create VPN connections to the appliance for the VPCs. Deploy instances in the new VPC in the Local Zone.
C- Create a new subnet in the Local Zone. Deploy a third-party appliance in the VPC with interfaces in each subnet. Configure the new subnet to route the Local Zone through the appliance. Deploy instances in the new subnet.
D- Create a new subnet in the Local Zone. Configure the new subnet to use a CIDR block that is within the VPC's CIDR block. Deploy instances in the new subnet in the Local Zone.

Answer:

D

Question 5

Question Type: MultipleChoice

A company recently implemented a security policy that prohibits developers from launching VPC network infrastructure. The policy states that any time a NAT gateway is launched in a VPC, the company's network security team must immediately receive an alert to terminate the NAT gateway. The network security team needs to implement a solution that can be deployed across AWS accounts with the least possible administrative overhead. The solution also must provide the network security team with a simple way to view compliance history.

Which solution will meet these requirements?

Options:

A- Develop a script that programmatically checks for NAT gateways in an AWS account, sends an email alert, and terminates the NAT gateway if a NAT gateway is detected. Deploy the script on an Amazon EC2 instance in each account. Use a cron job to run the script every 5 minutes. Log the results of the checks to an Amazon RDS for MySQL database.
B- Create an AWS Lambda function that programmatically checks for NAT gateways in an AWS account, sends an email alert, and terminates the NAT gateway if a NAT gateway is detected. Deploy the Lambda function to each account by using AWS Serverless Application Model (AWS SAM) templates. Store the results of the checks on an Amazon OpenSearch Service cluster in each account.
C- Enable Amazon GuardDuty. Create an Amazon EventBridge rule for the Behavior:EC2/NATGatewayCreation GuardDuty finding type. Configure the rule to invoke an AWS Step Functions state machine to send an email alert and terminate a NAT gateway if a NAT gateway is detected. Store the runtime log as a text file in an Amazon S3 bucket.
D- Create a custom AWS Config rule that checks for NAT gateways in an AWS account. Configure the AWS Config rule to perform an AWS Systems Manager Automation remediation action to send an email alert and terminate the NAT gateway if a NAT gateway is detected. Deploy the AWS Config rule and the Systems Manager runbooks to each account by using AWS CloudFormation StackSets

Answer:

D

Question 6

Question Type: MultipleChoice

An online retail company is running a web application in the us-west-2 Region and serves consumers in the United States. The company plans to expand across several countries in Europe and wants to provide low latency for all its users.

The application needs to identify the users' IP addresses and provide localized content based on the users' geographic location. The application uses HTTP GET and POST methods for its functionality. The company also needs to develop a failover mechanism that works for GET and POST methods and is based on health checks. The failover must occur in less than 1 minute for all clients.

Which solution will meet these requirements?

Options:

A- Configure a Network Load Balancer (NLB) for the application in each environment In the new AWS Regions. Create an AWS Global Accelerator accelerator that has endpoint groups that point to the NLBs in each Region.
B- Configure an Application Load Balancer (ALB) for the application in each environment in the new AWS Regions. Create an AWS Global Accelerator accelerator that has endpoint groups that

point to the ALBs in each Region.
C- Configure an Application Load Balancer (ALB) for the application in each environment in the new AWS Regions. Create Amazon Route 53 public hosted zones that have failover routing policies.
D- Configure a Network Load Balancer (NLB) for the application in each environment in the new AWS Regions. Create an Amazon CloudFront distribution. Configure an origin group with origin failover options.

Answer:

B

Question 7

Question Type: MultipleChoice

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

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

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

Options:

A- Create a VPN connection over the Direct Connect connection by using the on-premises firewall. Use the firewall to block all traffic from on premises to AWS. Allow a stateful connection from the EC2 instances to initiate the requests.
B- Configure the on-premises firewall to filter all requests from the on-premises network to the EC2 instances. Allow a stateful connection if the EC2 instances in the VPC initiate the traffic.
C- Deploy a NAT gateway into a private subnet in the VPC where the EC2 instances are deployed. Specify the NAT gateway type as private. Configure the on-premises firewall to allow connections from the IP address that is assigned to the NAT gateway.
D- Deploy a NAT instance into a private subnet in the VPC where the EC2 instances are deployed. Configure the on-premises firewall to allow connections from the IP address that is assigned to the NAT instance.

Answer:

C

Question 8

Question Type: MultipleChoice

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

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

Which solution will meet these requirements?

Options:

A- Configure multiple dynamic BGP-based Site-to-Site VPN connections to the transit gateway Configure equal-cost multi-path routing
B- Configure multiple static routing-based Site-to-Site VPN connections to the transit gateway. Configure equal-cost multi-path routing
C- Configure a new Site-to-Site VPN connection to the transit gateway Enable acceleration for the Site-to-Site VPN connection.
D- Configure a software appliance-based VPN connection over the internet from the on-premises firewall to an EC2 instance that has a large instance size and networking capabilities.

Answer:

A

Question 9

Question Type: MultipleChoice

A company runs a workload in a single VPC on AWS. The company's architecture contains several interface VPC endpoints for AWS services, including Amazon CloudWatch Logs and AWS Key Management Service (AWS KMS). The endpoints are configured to use a shared security group. The security group is not used for any other workloads or resources.

After a security review of the environment, the company determined that the shared security group is more permissive than necessary. The company wants to make the rules associated with the security group more restrictive. The changes to the security group rules must not prevent the resources in the VPC from using AWS services through interface VPC endpoints. The changes must prevent unnecessary access.

The security group currently uses the following rules:

* Inbound - Rule 1

Protocol: TCP

Port: 443

Source: 0.0.0.0/0

* Inbound - Rule 2

Protocol: TCP

Port: 443

Source: VPC CIDR

* Outbound - Rule 1

Protocol: All

Port: All

Destination: 0.0.0.0/0

Which rule or rules should the company remove to meet with these requirements?

Options:

A- Outbound - Rule 2
B- Inbound - Rule 1 and Outbound - Rule 1
C- Inbound - Rule 2 and Outbound - Rule 1
D- Outbound - Rule 1

Answer:

B

Explanation:

Inbound Rule 1 (Allow TCP 443 from 0.0.0.0/0): This rule allows all sources, including the public internet, to access the interface VPC endpoints. Since interface VPC endpoints are used within the VPC for communication with AWS services, this rule is unnecessarily permissive. Removing this rule enhances security while still allowing communication within the VPC using Rule 2 (TCP 443 from the VPC CIDR).

Outbound Rule 1 (Allow All Protocols, All Ports to 0.0.0.0/0): This rule is overly permissive and unnecessary for interface VPC endpoints, as traffic destined for AWS services through these

endpoints does not need unrestricted outbound access. Removing this rule ensures that outbound traffic is limited to what is required for communication with the AWS services through the interface endpoints.

P2P exams