

# Amazon-Web-Services

## Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam



**NEW QUESTION 1**

A company has an AWS Site-to-Site VPN connection between its existing VPC and on-premises network. The default DHCP options set is associated with the VPC. The company has an application that is running on an Amazon Linux 2 Amazon EC2 instance in the VPC. The application must retrieve an Amazon RDS database secret that is stored in AWS Secrets Manager through a private VPC endpoint. An on-premises application provides internal RESTful API service that can be reached by URL (<https://api.example.internal>). Two on-premises Windows DNS servers provide internal DNS resolution. The application on the EC2 instance needs to call the internal API service that is deployed in the on-premises environment. When the application on the EC2 instance attempts to call the internal API service by referring to the hostname that is assigned to the service, the call fails. When a network engineer tests the API service call from the same EC2 instance by using the API service's IP address, the call is successful. What should the network engineer do to resolve this issue and prevent the same problem from affecting other resources in the VPC?

- A. Create a new DHCP options set that specifies the on-premises Windows DNS server
- B. Associate the new DHCP options set with the existing VP
- C. Reboot the Amazon Linux 2 EC2 instance.
- D. Create an Amazon Route 53 Resolver rul
- E. Associate the rule with the VP
- F. Configure the rule to forward DNS queries to the on-premises Windows DNS servers if the domain name matches example.internal.
- G. Modify the local host file in the Amazon Linux 2 EC2 instance in the VPMap the service domain name (api.example.internal) to the IP address of the internal API service.
- H. Modify the local /etc/resolv.conf file in the Amazon Linux 2 EC2 instance in the VP
- I. Change the IP addresses of the name servers in the file to the IP addresses of the company's on-premises Windows DNS servers.

**Answer:** B

**Explanation:**

Creating an Amazon Route 53 Resolver rule and associating it with the VPC would enable forwarding of DNS queries for a specified domain name (example.internal) to a specified IP address (the on-premises Windows DNS servers)<sup>3</sup>. This would allow EC2 instances in the VPC to resolve the internal API service by using its hostname. Configuring the rule to forward DNS queries only if the domain name matches example.internal would also allow EC2 instances to use the Amazon Route 53 Resolver server for other DNS queries, such as those for AWS services through private VPC endpoints<sup>2</sup>.

**NEW QUESTION 2**

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

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

Which solution will meet these requirements?

- A. Configure the ALB in a private subnet of the VP
- B. Attach an internet gateway without adding routes in the subnet route tables to point to the internet gatewa
- C. Configure the accelerator with endpoint groups that include the ALB endpoint
- D. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- E. Configure the ALB in a private subnet of the VP
- F. Configure the accelerator with endpoint groups that include the ALB endpoint
- G. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- H. Configure the ALB in a public subnet of the VPAttach an internet gatewa
- I. Add routes in the subnet route tables to point to the internet gatewa
- J. Configure the accelerator with endpoint groups that include the ALB endpoint
- K. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.
- L. Configure the ALB in a private subnet of the VP
- M. Attach an internet gatewa
- N. Add routes in the subnet route tables to point to the internet gatewa
- O. Configure the accelerator with endpoint groups that include the ALB endpoint
- P. Configure the ALB's security group to only allow inbound trafficfrom the accelerator's IP addresses on the ALB listener port.

**Answer:** A

**Explanation:**

Please read the below link typically describing ELB integration with AWS Global accelator (and the last line of the extract) - <https://docs.aws.amazon.com/global-accelerator/latest/dg/secure-vpc-connections.html> "When you add an internal Application Load Balancer or an Amazon EC2 instance endpoint in AWS Global Accelerator, you enable internet traffic to flow directly to and from the endpoint in Virtual Private Clouds (VPCs) by targeting it in a private subnet. The VPC that contains the load balancer or EC2 instance must have an internet gateway attached to it, to indicate that the VPC accepts internet traffic. However, you don't need public IP addresses on the load balancer or EC2 instance. You also don't need an associated internet gateway route for the subnet."

**NEW QUESTION 3**

A network engineer needs to standardize a company's approach to centralizing and managing interface VPC endpoints for private communication with AWS services. The company uses AWS Transit Gateway for inter-VPC connectivity between AWS accounts through a hub-and-spoke model. The company's network services team must manage all Amazon Route 53 zones and interface endpoints within a shared services AWS account. The company wants to use thiscentralized model to provide AWS resources with access to AWS Key Management Service (AWS KMS) without sending traffic over the public internet. What should the network engineer do to meet these requirements?

- A. In the shared services account, create an interface endpoint for AWS KM
- B. Modify the interface endpoint by disabling the private DNS nam
- C. Create a private hosted zone in the shared services account with an alias record that points to the interface endpoint
- D. Associate the private hosted zone with the spoke VPCs in each AWS account.
- E. In the shared services account, create an interface endpoint for AWS KM
- F. Modify the interface endpoint by disabling the private DNS nam
- G. Create a private hosted zone in each spoke AWS account with an alias record that points to the interface endpoint

- H. Associate each private hosted zone with the shared services AWS account.
- I. In each spoke AWS account, create an interface endpoint for AWS KM
- J. Modify each interface endpoint by disabling the private DNS nam
- K. Create a private hosted zone in each spoke AWS account with an alias record that points to each interface endpoint
- L. Associate each private hosted zone with the shared services AWS account.
- M. In each spoke AWS account, create an interface endpoint for AWS KM
- N. Modify each interface endpoint by disabling the private DNS nam
- O. Create a private hosted zone in the shared services account with an alias record that points to each interface endpoint
- P. Associate the private hosted zone with the spoke VPCs in each AWS account.

**Answer:** A

#### NEW QUESTION 4

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

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

**Answer:** D

#### NEW QUESTION 5

A company has an AWS Direct Connect connection between its on-premises data center in the United States (US) and workloads in the us-east-1 Region. The connection uses a transit VIF to connect the data center to a transit gateway in us-east-1.

The company is opening a new office in Europe with a new on-premises data center in England. A Direct Connect connection will connect the new data center with some workloads that are running in a single VPC in the eu-west-2 Region. The company needs to connect the US data center and us-east-1 with the Europe data center and eu-west-2. A network engineer must establish full connectivity between the data centers and Regions with the lowest possible latency. How should the network engineer design the network architecture to meet these requirements?

- A. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VI
- B. Associate the transit gateway in us-east-1 with the same Direct Connect gatewa
- C. Enable SiteLink for the transit VIF and the private VIF.
- D. Connect the VPC in eu-west-2 to a new transit gatewa
- E. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VI
- F. Associate the transit gateway in us-east-1 with the same Direct Connect gatewa
- G. Enable SiteLink for both transit VIF
- H. Peer the two transit gateways.
- I. Connect the VPC in eu-west-2 to a new transit gatewa
- J. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VI
- K. Create a new Direct Connect gatewa
- L. Associate the transit gateway in us-east-1 with the new Direct Connect gatewa
- M. Enable SiteLink for both transit VIF
- N. Peer the two transit gateways.
- O. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VI
- P. Create a new Direct Connect gatewa
- Q. Associate the transit gateway in us-east-1 with the new Direct Connect gatewa
- R. Enable SiteLink for the transit VIF and the private VIF.

**Answer:** C

#### NEW QUESTION 6

An insurance company is planning the migration of workloads from its on-premises data center to the AWS Cloud. The company requires end-to-end domain name resolution. Bi-directional DNS resolution between AWS and the existing on-premises environments must be established. The workloads will be migrated into multiple VPCs. The workloads also have dependencies on each other, and not all the workloads will be migrated at the same time.

Which solution meets these requirements?

- A. Configure a private hosted zone for each application VPC, and create the requisite record
- B. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VP
- C. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolve
- D. Associate the application VPC private hosted zones with the egress VPC, and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manage
- E. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inboundendpoints.
- F. Configure a public hosted zone for each application VPC, and create the requisite record
- G. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VP
- H. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolve
- I. Associate the application VPC private hosted zones with the egress VP
- J. and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manage
- K. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints.
- L. Configure a private hosted zone for each application VPC, and create the requisite record
- M. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPDefine Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolve
- N. Associate the application VPC private hosted zones with the egress VPand s

**Answer:** A

**Explanation:**

Creating a private hosted zone for each application VPC and creating the requisite records would enable end-to-end domain name resolution for the resources. Creating a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC would enable bi-directional DNS resolution between AWS and the existing on-premises environments. Defining Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver would enable DNS queries from AWS resources to on-premises resources. Associating the application VPC private hosted zones with the egress VPC and sharing the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager would enable DNS queries among different VPCs and accounts. Configuring the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints would enable DNS queries from on-premises resources to AWS resources1.

**NEW QUESTION 7**

A network engineer needs to set up an Amazon EC2 Auto Scaling group to run a Linux-based network appliance in a highly available architecture. The network engineer is configuring the new launch template for the Auto Scaling group.

In addition to the primary network interface the network appliance requires a second network interface that will be used exclusively by the application to exchange traffic with hosts over the internet. The company has set up a Bring Your Own IP (BYOIP) pool that includes an Elastic IP address that should be used as the public IP address for the second network interface.

How can the network engineer implement the required architecture?

- A. Configure the two network interfaces in the launch template
- B. Define the primary network interface to be created in one of the private subnet
- C. For the second network interface, select one of the public subnet
- D. Choose the BYOIP pool ID as the source of public IP addresses.
- E. Configure the primary network interface in a private subnet in the launch template
- F. Use the user data option to run a cloud-init script after boot to attach the second network interface from a subnet with auto-assign public IP addressing enabled.
- G. Create an AWS Lambda function to run as a lifecycle hook of the Auto Scaling group when an instance is launched
- H. In the Lambda function, assign a network interface to an AWS Global Accelerator endpoint.
- I. During creation of the Auto Scaling group, select subnets for the primary network interface
- J. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

**Answer: D**

**Explanation:**

During creation of the Auto Scaling group, select subnets for the primary network interface. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

This solution meets all of the requirements stated in the question. The primary network interface can be configured in a private subnet during creation of the Auto Scaling group. The user data option can be used to run a cloud-init script that will allocate a second network interface and associate an Elastic IP address from the BYOIP pool with it.

**NEW QUESTION 8**

Your organization has a newly installed 1-Gbps AWS Direct Connect connection. You order the cross-connect from the Direct Connect location provider to the port on your router in the same facility. To enable the use of your first virtual interface, your router must be configured appropriately.

What are the minimum requirements for your router?

- A. 1-Gbps Multi Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
- B. 1-Gbps Single Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
- C. IPsec Parameters, Pre-Shared key, Peer IP Address, BGP Session with MD5
- D. BGP Session with MD5, 802.1Q VLAN, Route-Map, Prefix List, IPsec encrypted GRE Tunnel

**Answer: B**

**NEW QUESTION 9**

A software-as-a-service (SaaS) provider hosts its solution on Amazon EC2 instances within a VPC in the AWS Cloud. All of the provider's customers also have their environments in the AWS Cloud.

A recent design meeting revealed that the customers have IP address overlap with the provider's AWS deployment. The customers have stated that they will not share their internal IP addresses and that they do not want to connect to the provider's SaaS service over the internet.

Which combination of steps is part of a solution that meets these requirements? (Choose two.)

- A. Deploy the SaaS service endpoint behind a Network Load Balancer.
- B. Configure an endpoint service, and grant the customers permission to create a connection to the endpoint service.
- C. Deploy the SaaS service endpoint behind an Application Load Balancer.
- D. Configure a VPC peering connection to the customer VPC
- E. Route traffic through NAT gateways.
- F. Deploy an AWS Transit Gateway, and connect the SaaS VPC to it
- G. Share the transit gateway with the customer
- H. Configure routing on the transit gateway.

**Answer: AB**

**Explanation:**

NLB for creating the private link which solves the overlapping IP address issue and the SaaS service endpoint behind it. (the SaaS endpoint could be an ALB)  
<https://aws.amazon.com/about-aws/whats-new/2021/09/application-load-balancer-aws-privatelink-static-ip>

**NEW QUESTION 10**

A customer has set up multiple VPCs for Dev, Test, Prod, and Management. You need to set up AWS Direct Connect to enable data flow from on-premises to each VPC. The customer has monitoring software running in the Management VPC that collects metrics from the instances in all the other VPCs. Due to budget requirements, data transfer charges should be kept at minimum.

Which design should be recommended?

- A. Create a total of four private VIFs, one for each VPC owned by the customer, and route traffic between VPCs using the Direct Connect link.
- B. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs.



- C. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs, enable source/destination NAT in the Management VPC.
- D. Create a total of four private VIFs, and enable VPC peering between all VPCs.

**Answer:** D

**Explanation:**

- creating VPC peering is free of charge - traffic costs ~0.01€/GB for VPC peering (IN + OUT) and ~0.02€/GB for direct connect (OUT only). As the communication involved in monitoring will never have IN == OUT, then  $0.01 * (IN + OUT)$  will always be lower than  $0.02 * OUT$ , ergo VPC peering will be cheaper

**NEW QUESTION 10**

You deploy an Amazon EC2 instance that runs a web server into a subnet in a VPC. An Internet gateway is attached, and the main route table has a default route (0.0.0.0/0) configured with a target of the Internet gateway.

The instance has a security group configured to allow as follows:

- Protocol: TCP
- Port: 80 inbound, nothing outbound

The Network ACL for the subnet is configured to allow as follows:

- Protocol: TCP
- Port: 80 inbound, nothing outbound

When you try to browse to the web server, you receive no response. Which additional step should you take to receive a successful response?

- A. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 80
- B. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 1024-65535
- C. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 80
- D. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 1024-65535

**Answer:** D

**Explanation:**

To enable the connection to a service running on an instance, the associated network ACL must allow both inbound traffic on the port that the service is listening on as well as allow outbound traffic from ephemeral ports. When a client connects to a server, a random port from the ephemeral port range (1024-65535) becomes the client's source port. The designated ephemeral port then becomes the destination port for return traffic from the service, so outbound traffic from the ephemeral port must be allowed in the network ACL. <https://aws.amazon.com/premiumsupport/knowledge-center/resolve-connection-sg-acl-inbound/>

**NEW QUESTION 15**

A company uses a 1 Gbps AWS Direct Connect connection to connect its AWS environment to its on-premises data center. The connection provides employees with access to an application VPC that is hosted on AWS. Many remote employees use a company-provided VPN to connect to the data center. These employees are reporting slowness when they access the application during business hours. On-premises users have started to report similar slowness while they are in the office.

The company plans to build an additional application on AWS. On-site and remote employees will use the additional application. After the deployment of this additional application, the company will need 20% more bandwidth than the company currently uses. With the increased usage, the company wants to add resiliency to the AWS connectivity. A network engineer must review the current implementation and must make improvements within a limited budget.

What should the network engineer do to meet these requirements MOST cost-effectively?

- A. Set up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional traffic load from remote employees and the additional application
- B. Create a link aggregation group (LAG).
- C. Deploy an AWS Site-to-Site VPN connection to the application VP
- D. Configure the on-premises routing for the remote employees to connect to the Site-to-Site VPN connection.
- E. Deploy Amazon Workspaces into the application VPC and instruct the remote employees to connect to Workspaces.
- F. Replace the existing 1 Gbps Direct Connect connection with two new 2 Gbps Direct Connect hosted connections
- G. Create an AWS Client VPN endpoint in the application VP
- H. Instruct the remote employees to connect to the Client VPN endpoint.

**Answer:** A

**Explanation:**

Setting up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional traffic load from remote employees and the additional application would provide more bandwidth and lower latency than a VPN connection over the public internet<sup>1</sup>. Creating a link aggregation group (LAG) with the existing and new Direct Connect connections would provide resiliency and redundancy for the AWS connectivity<sup>2</sup>.

**NEW QUESTION 20**

An IoT company sells hardware sensor modules that periodically send out temperature, humidity, pressure, and location data through the MQTT messaging protocol. The hardware sensor modules send this data to the company's on-premises MQTT brokers that run on Linux servers behind a load balancer. The hardware sensor modules have been hardcoded with public IP addresses to reach the brokers.

The company is growing and is acquiring customers across the world. The existing solution can no longer scale and is introducing additional latency because of the company's global presence. As a result, the company decides to migrate its entire infrastructure from on-premises to the AWS Cloud. The company needs to migrate without reconfiguring the hardware sensor modules that are already deployed across the world. The solution also must minimize latency.

The company migrates the MQTT brokers to run on Amazon EC2 instances. What should the company do next to meet these requirements?

- A. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listener
- B. Use Bring Your Own IP (BYOIP) from the on-premises network with the NLB.
- C. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listener
- D. Create an AWS Global Accelerator accelerator in front of the NLB. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator.
- E. Place the EC2 instances behind an Application Load Balancer (ALB). Configure TCP listener
- F. Create an AWS Global Accelerator accelerator in front of the ALB
- G. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator
- H. Place the EC2 instances behind an Amazon CloudFront distribution
- I. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

**Answer:** B

#### NEW QUESTION 23

A global company runs business applications in the us-east-1 Region inside a VPC. One of the company's regional offices in London uses a virtual private gateway for an AWS Site-to-Site VPN connection to the VPC. The company has configured a transit gateway and has set up peering between the VPC and other VPCs that various departments in the company use.

Employees at the London office are experiencing latency issues when they connect to the business applications.

What should a network engineer do to reduce this latency?

- A. Create a new Site-to-Site VPN connection
- B. Set the transit gateway as the target gateway
- C. Enable acceleration on the new Site-to-Site VPN connection
- D. Update the VPN device in the London office with the new connection details.
- E. Modify the existing Site-to-Site VPN connection by setting the transit gateway as the target gateway. Enable acceleration on the existing Site-to-Site VPN connection.
- F. Create a new transit gateway in the eu-west-2 (London) Region
- G. Peer the new transit gateway with the existing transit gateway
- H. Modify the existing Site-to-Site VPN connection by setting the new transit gateway as the target gateway.
- I. Create a new AWS Global Accelerator standard accelerator that has an endpoint of the Site-to-Site VPN connection
- J. Update the VPN device in the London office with the new connection details.

**Answer:** A

#### Explanation:

Enabling acceleration for a Site-to-Site VPN connection uses AWS Global Accelerator to route traffic from the on-premises network to an AWS edge location that is closest to the customer gateway device<sup>1</sup>. AWS Global Accelerator optimizes the network path, using the congestion-free AWS global network to route traffic to the endpoint that provides the best application performance<sup>2</sup>. Setting the transit gateway as the target gateway enables connectivity between the on-premises network and multiple VPCs that are attached to the transit gateway<sup>3</sup>.

#### NEW QUESTION 28

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

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

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

- A. Create a central transit gateway
- B. Create a VPC attachment to each application VPC
- C. Provide full mesh connectivity between all the VPCs by using the transit gateway.
- D. Create VPC peering connections between the central shared services VPC and each application VPC in each business unit's AWS account.
- E. Create VPC endpoint services powered by AWS PrivateLink in the central shared services VPC. Create VPC endpoints in each application VPC.
- F. Create a central transit VPC with a VPN appliance from AWS Marketplace
- G. Create a VPN attachment from each VPC to the transit VPC
- H. Provide full mesh connectivity among all the VPCs.

**Answer:** C

#### Explanation:

Option C provides a secure and scalable solution using VPC endpoint services powered by AWS PrivateLink. AWS PrivateLink enables private connectivity between VPCs and services without exposing the data to the public internet or using a VPN connection. By creating VPC endpoints in each application VPC, the company can securely access the central shared services VPC without the need for complex network configurations. Furthermore, PrivateLink supports cross-account connectivity, which makes it a scalable solution as more business units consume data from the central shared services VPC in the future.

#### NEW QUESTION 29

A company is using Amazon Route 53 Resolver DNS Firewall in a VPC to block all domains except domains that are on an approved list. The company is concerned that if DNS Firewall is unresponsive, resources in the VPC might be affected if the network cannot resolve any DNS queries. To maintain application service level agreements, the company needs DNS queries to continue to resolve even if Route 53 Resolver does not receive a response from DNS Firewall.

Which change should a network engineer implement to meet these requirements?

- A. Update the DNS Firewall VPC configuration to disable fail open for the VPC.
- B. Update the DNS Firewall VPC configuration to enable fail open for the VPC.
- C. Create a new DHCP options set with parameter `dns_firewall_fail_open=false`
- D. Associate the new DHCP options set with the VPC.
- E. Create a new DHCP options set with parameter `dns_firewall_fail_open=true`
- F. Associate the new DHCP options set with the VPC.

**Answer:** B

#### NEW QUESTION 32

An AWS CloudFormation template is being used to create a VPC peering connection between two existing operational VPCs, each belonging to a different AWS account. All necessary components in the 'Remote' (receiving) account are already in place.

The template below creates the VPC peering connection in the Originating account. It contains these components:

AWSTemplateFormation Version: 2010-09-09 Parameters:

Originating VPCId: Type: String RemoteVPCId: Type: String

RemoteVPCAccountId: Type: String Resources:

newVPCPeeringConnection:

Type: 'AWS::EC2::VPCPeeringConnection'

Properties:

VpcId: !Ref OriginatingVPCId PeerVpcId: !Ref RemoteVPCId PeerOwnerId: !Ref RemoteVPCAccountId

Which additional AWS CloudFormation components are necessary in the Originating account to create an operational cross-account VPC peering connection with AWS CloudFormation? (Select two.)

- A. Resources:NewEC2SecurityGroup:Type: AWS::EC2::SecurityGroup
- B. Resources:NetworkInterfaceToRemoteVPC:Type: "AWS::EC2::NetworkInterface"
- C. Resources:newEC2Route:Type: AWS::EC2::Route
- D. Resources:VPCGatewayToRemoteVPC:Type: "AWS::EC2::VPCGatewayAttachment"
- E. Resources:newVPCPeeringConnection:Type: 'AWS::EC2VPCPeeringConnection'PeerRoleArn: !Ref PeerRoleArn

**Answer:** CE

**Explanation:**

[https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/AWS\\_EC2.html](https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/AWS_EC2.html)

#### NEW QUESTION 34

A company is deploying an application. The application is implemented in a series of containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use the Fargate launch type for its tasks. The containers will run workloads that require connectivity initiated over an SSL connection. Traffic must be able to flow to the application from other AWS accounts over private connectivity. The application must scale in a manageable way as more consumers use the application.

Which solution will meet these requirements?

- A. Choose a Gateway Load Balancer (GLB) as the type of load balancer for the ECS service
- B. Create a lifecycle hook to add new tasks to the target group from Amazon ECS as required to handle scaling
- C. Specify the GLB in the service definition
- D. Create a VPC peer for external AWS account
- E. Update the route tables so that the AWS accounts can reach the GLB.
- F. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service
- G. Create path-based routing rules to allow the application to target the containers that are registered in the target group
- H. Specify the ALB in the service definition
- I. Create a VPC endpoint service for the ALB. Share the VPC endpoint service with other AWS accounts.
- J. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service
- K. Create path-based routing rules to allow the application to target the containers that are registered in the target group
- L. Specify the ALB in the service definition
- M. Create a VPC peer for the external AWS account
- N. Update the route tables so that the AWS accounts can reach the ALB.
- O. Choose a Network Load Balancer (NLB) as the type of load balancer for the ECS service
- P. Specify the NLB in the service definition
- Q. Create a VPC endpoint service for the NLB
- R. Share the VPC endpoint service with other AWS accounts.

**Answer:** D

#### NEW QUESTION 35

A company plans to deploy a two-tier web application to a new VPC in a single AWS Region. The company has configured the VPC with an internet gateway and four subnets. Two of the subnets are public and have default routes that point to the internet gateway. Two of the subnets are private and share a route table that does not have a default route.

The application will run on a set of Amazon EC2 instances that will be deployed behind an external Application Load Balancer. The EC2 instances must not be directly accessible from the internet. The application will use an Amazon S3 bucket in the same Region to store data. The application will invoke S3 GET API operations and S3 PUT API operations from the EC2 instances. A network engineer must design a VPC architecture that minimizes data transfer cost.

Which solution will meet these requirements?

- A. Deploy the EC2 instances in the public subnet
- B. Create an S3 interface endpoint in the VPC
- C. Modify the application configuration to use the S3 endpoint-specific DNS hostname.
- D. Deploy the EC2 instances in the private subnet
- E. Create a NAT gateway in the VPC
- F. Create default routes in the private subnets to the NAT gateway
- G. Connect to Amazon S3 by using the NAT gateway.
- H. Deploy the EC2 instances in the private subnet
- I. Create an S3 gateway endpoint in the VPC. Specify the route table of the private subnets during endpoint creation to create routes to Amazon S3.
- J. Deploy the EC2 instances in the private subnet
- K. Create an S3 interface endpoint in the VPC
- L. Modify the application configuration to use the S3 endpoint-specific DNS hostname.

**Answer:** C

**Explanation:**

Option C is the optimal solution as it involves deploying the EC2 instances in the private subnets, which provides additional security benefits. Additionally, creating an S3 gateway endpoint in the VPC will enable the EC2 instances to communicate with Amazon S3 directly, without incurring data transfer costs. This is because the S3 gateway endpoint uses Amazon's private network to transfer data between the VPC and S3, which is not charged for data transfer. Furthermore, specifying the route table of the private subnets during endpoint creation will create routes to Amazon S3, which is required for the EC2 instances to communicate with S3.

#### NEW QUESTION 39

A company has created three VPCs: a production VPC, a nonproduction VPC, and a shared services VPC. The production VPC and the nonproduction VPC must each have communication with the shared services VPC. There must be no communication between the production VPC and the nonproduction VPC. A transit gateway is deployed to facilitate communication between VPCs.

Which route table configurations on the transit gateway will meet these requirements?

- A. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for only the shared services VPC



- B. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.
- C. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for each VP
- D. Create an additional route table with only the shared services VPC attachment associated with propagated routes from each VPC.
- E. Configure a route table with all the VPC attachments associated with propagated routes for only the shared services VPC
- F. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.
- G. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes disable
- G. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.

**Answer:** A

#### NEW QUESTION 41

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

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

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

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

**Answer:** BDF

#### NEW QUESTION 43

A company is deploying a non-web application on an AWS load balancer. All targets are servers located on-premises that can be accessed by using AWS Direct Connect. The company wants to ensure that the source IP addresses of clients connecting to the application are passed all the way to the end server.

How can this requirement be achieved?

- A. Use a Network Load Balancer to automatically preserve the source IP address.
- B. Use a Network Load Balancer and enable the X-Forwarded-For attribute.
- C. Use a Network Load Balancer and enable the ProxyProtocol v2 attribute.
- D. Use an Application Load Balancer to automatically preserve the source IP address in the X-Forwarded-For header.

**Answer:** C

#### Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#proxy-protocol>

#### NEW QUESTION 45

A company's development team has created a new product recommendation web service. The web service is hosted in a VPC with a CIDR block of 192.168.224.0/19. The company has deployed the web service on Amazon EC2 instances and has configured an Auto Scaling group as the target of a Network Load Balancer (NLB).

The company wants to perform testing to determine whether users who receive product recommendations spend more money than users who do not receive product recommendations. The company has a big sales event in 5 days and needs to integrate its existing production environment with the recommendation engine by then. The existing production environment is hosted in a VPC with a CIDR block of 192.168.128.0/17.

A network engineer must integrate the systems by designing a solution that results in the least possible disruption to the existing environments.

Which solution will meet these requirements?

- A. Create a VPC peering connection between the web service VPC and the existing production VP
- B. Add a routing rule to the appropriate route table to allow data to flow to 192.168.224.0/19 from the existing production environment and to flow to 192.168.128.0/17 from the web service environmen
- C. Configure the relevant security groups and ACLs to allow the systems to communicate.
- D. Ask the development team of the web service to redeploy the web service into the production VPC and integrate the systems there.
- E. Create a VPC endpoint servic
- F. Associate the VPC endpoint service with the NLB for the web service. Create an interface VPC endpoint for the web service in the existing production VPC.
- G. Create a transit gateway in the existing production environmen
- H. Create attachments to the production VPC and the web service VP
- I. Configure appropriate routing rules in the transit gateway and VPC route tables for 192.168.224.0/19 and 192.168.128.0/17. Configure the relevant security groups and ACLs to allow the systems to communicate.

**Answer:** C

#### NEW QUESTION 50

A company hosts an application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company recently experienced a network security breach. A network engineer must collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application.

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

- A. Configure the ALB to store logs in an Amazon S3 bucke
- B. Download the files from Amazon S3, and use a spreadsheet application to analyze the logs.
- C. Configure the ALB to push logs to Amazon Kinesis Data Stream
- D. Use Amazon Kinesis Data Analytics to analyze the logs.
- E. Configure Amazon Kinesis Data Streams to stream data from the ALB to Amazon OpenSearch Service (Amazon Elasticsearch Service). Use search operations



in Amazon OpenSearch Service (Amazon Elasticsearch Service) to analyze the data.  
F. Configure the ALB to store logs in an Amazon S3 bucket  
G. Use Amazon Athena to analyze the logs in Amazon S3.

**Answer: D**

**Explanation:**

The most operationally efficient solution to collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application would be to configure the ALB to store logs in an Amazon S3 bucket and use Amazon Athena to analyze the logs in Amazon S3 (Option D). This solution allows for quick and easy analysis of log data without requiring manual download or manipulation of log files.

**NEW QUESTION 53**

A company has deployed a critical application on a fleet of Amazon EC2 instances behind an Application Load Balancer. The application must always be reachable on port 443 from the public internet. The application recently had an outage that resulted from an incorrect change to the EC2 security group. A network engineer needs to automate a way to verify the network connectivity between the public internet and the EC2 instances whenever a change is made to the security group. The solution also must notify the network engineer when the change affects the connection. Which solution will meet these requirements?

- A. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture REJECT traffic on port 443. Publish the flow log records to a log group in Amazon CloudWatch Log
- B. Create a CloudWatch Logs metric filter for the log group for rejected traffic
- C. Create an alarm to notify the network engineer.
- D. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture all traffic on port 443. Publish the flow log records to a log group in Amazon CloudWatch Log
- E. Create a CloudWatch Logs metric filter for the log group for all traffic
- F. Create an alarm to notify the network engineer
- G. Create a VPC Reachability Analyzer path on port 443. Specify the security group as the source
- H. Specify the EC2 instances as the destination
- I. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection
- J. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail
- K. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs
- L. Specify the EC2 instances as the destination
- M. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection
- N. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail
- O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs.

**Answer: C**

**NEW QUESTION 57**

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

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

**Answer: C**

**Explanation:**

Using AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between VPCs and transit gateways<sup>1</sup>. Verifying that the VPC route tables are correct would enable identification of routing issues within a VPC. Using VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC would enable identification of traffic filtering issues within a VPC<sup>2</sup>. Additionally, using VPC Reachability Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between transit gateways in different Regions. VPC Reachability Analyzer is a configuration analysis tool that enables connectivity testing between a source resource and a destination resource in your VPCs.

**NEW QUESTION 62**

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

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

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

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

- H. Deploy a new pair of 10 GB Direct Connect connections with MACse
- I. Deploy a new public VIF on the new Direct Connect connection
- J. Deploy two AWS Site-to-Site VPN connections on top of the new public VI
- K. Reroute traffic from the existing private VIF to the new Site-to-Site connection
- L. Decommission the original Direct Connect connections.

**Answer:** C

#### NEW QUESTION 63

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

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

What action will resolve the availability problem?

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

**Answer:** B

#### NEW QUESTION 67

A company's network engineer is designing an active-passive connection to AWS from two on-premises data centers. The company has set up AWS Direct Connect connections between the on-premises data centers and AWS. From each location, the company is using a transit VIF that connects to a Direct Connect gateway that is associated with a transit gateway.

The network engineer must ensure that traffic from AWS to the data centers is routed first to the primary data center. The traffic should be routed to the failover data center only in the case of an outage.

Which solution will meet these requirements?

- A. Set the BGP community tag for all prefixes from the primary data center to 7224:7100. Set the BGP community tag for all prefixes from the failover data center to 7224:7300
- B. Set the BGP community tag for all prefixes from the primary data center to 7224:7300. Set the BGP community tag for all prefixes from the failover data center to 7224:7100
- C. Set the BGP community tag for all prefixes from the primary data center to 7224:9300. Set the BGP community tag for all prefixes from the failover data center to 7224:9100
- D. Set the BGP community tag for all prefixes from the primary data center to 7224:9100. Set the BGP community tag for all prefixes from the failover data center to 7224:9300

**Answer:** B

#### NEW QUESTION 71

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

Which solution will meet these requirements?

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

**Answer:** D

#### NEW QUESTION 76

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

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

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

Which solution will meet these requirements with the HIGHEST availability?

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

**Answer:** B

**Explanation:**

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

#### NEW QUESTION 78

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

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

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

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

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

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

**Answer:** AD

#### NEW QUESTION 80

A company has been using an outdated application layer protocol for communication among applications. The company decides not to use this protocol anymore and must migrate all applications to support a new protocol. The old protocol and the new protocol are TCP-based, but the protocols use different port numbers. After several months of work, the company has migrated dozens of applications that run on Amazon EC2 instances and in containers. The company believes that all the applications have been migrated, but the company wants to verify this belief. A network engineer needs to verify that no application is still using the old protocol.

Which solution will meet these requirements without causing any downtime?

- A. Use Amazon Inspector and its Network Reachability rules package
- B. Wait until the analysis has finished running to find out which EC2 instances are still listening to the old port.
- C. Enable Amazon GuardDuty
- D. Use the graphical visualizations to filter for traffic that uses the port of the old protocol
- E. Exclude all internet traffic to filter out occasions when the same port is used as an ephemeral port.
- F. Configure VPC flow logs to be delivered into an Amazon S3 bucket
- G. Use Amazon Athena to query the data and to filter for the port number that is used by the old protocol.
- H. Inspect all security groups that are assigned to the EC2 instances that host the application
- I. Remove the port of the old protocol if that port is in the list of allowed ports
- J. Verify that the applications are operating properly after the port is removed from the security groups.

**Answer:** C

**Explanation:**

Configuring VPC flow logs to be delivered into an Amazon S3 bucket would enable capture of information about the IP traffic going to and from network interfaces within the VPC. Using Amazon Athena to query the data and to filter for the port number that is used by the old protocol would enable identification of applications that are still using the old protocol.

#### NEW QUESTION 84



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

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 88**

A company has deployed a new web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Amazon EC2 Auto Scaling group. Enterprise customers from around the world will use the application. Employees of these enterprise customers will connect to the application over HTTPS from office locations.

The company must configure firewalls to allow outbound traffic to only approved IP addresses. The employees of the enterprise customers must be able to access the application with the least amount of latency.

Which change should a network engineer make in the infrastructure to meet these requirements?

- A. Create a new Network Load Balancer (NLB). Add the ALB as a target of the NLB.
- B. Create a new Amazon CloudFront distributio
- C. Set the ALB as the distribution's origin.
- D. Create a new accelerator in AWS Global Accelerato
- E. Add the ALB as an accelerator endpoint.
- F. Create a new Amazon Route 53 hosted zon
- G. Create a new record to route traffic to the ALB.

**Answer: B**

**Explanation:**

Amazon CloudFront is a content delivery network (CDN) that can speed up the delivery of static and dynamic web content, such as images, videos, and APIs<sup>2</sup>. CloudFront can also provide end-to-end encryption for HTTPS traffic by using SSL certificates from AWS Certificate Manager (ACM) or other sources<sup>2</sup>. CloudFront can also support session affinity (sticky sessions) with a load balancer-generated cookie or an application-based cookie policy<sup>2</sup>.

**NEW QUESTION 90**

A company has deployed an application in a VPC that uses a NAT gateway for outbound traffic to the internet. A network engineer notices a large quantity of suspicious network traffic that is traveling from the VPC over the internet to IP addresses that are included on a deny list. The network engineer must implement a solution to determine which AWS resources are generating the suspicious traffic. The solution must minimize cost and administrative overhead.

Which solution will meet these requirements?

- A. Launch an Amazon EC2 instance in the VP
- B. Use Traffic Mirroring by specifying the NAT gateway as the source and the EC2 instance as the destinatio
- C. Analyze the captured traffic by using open-source tools to identify the AWS resources that are generating the suspicious traffic.
- D. Use VPC flow log
- E. Launch a security information and event management (SIEM) solution in the VP
- F. Configure the SIEM solution to ingest the VPC flow log
- G. Run queries on the SIEM solution to identify the AWS resources that are generating the suspicious traffic.
- H. Use VPC flow log
- I. Publish the flow logs to a log group in Amazon CloudWatch Log
- J. Use CloudWatch Logs Insights to query the flow logs to identify the AWS resources that are generating the suspicious traffic.
- K. Configure the VPC to stream the network traffic directly to an Amazon Kinesis data strea
- L. Send the data from the Kinesis data stream to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Athena to query the data to identify the AWS resources that are generating the suspicious traffic.

**Answer: C**

**NEW QUESTION 93**

A bank built a new version of its banking application in AWS using containers that connect 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?

- 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 applicatio
- C. Route all traffic to the new application by using an Elastic Load Balancing (ELB) load balancer DN
- D. Define a user-agent-based rule on the backend servers to redirect earlier clients to the on-premises application.
- E. Use an Application Load Balancer for the new applicatio
- F. Register both the new and earlier applications as separate target groups and use path-based routing to route traffic based on the application version.
- G. Use an Application Load Balancer for the new applicatio
- H. Register both the new and earlier application backends as separate target group
- I. Use header-based routing to route traffic based on the application version.

**Answer: D**

#### NEW QUESTION 94

A company wants to improve visibility into its AWS environment. The AWS environment consists of multiple VPCs that are connected to a transit gateway. The transit gateway connects to an on-premises data center through an AWS Direct Connect gateway and a pair of redundant Direct Connect connections that use transit VIFs. The company must receive notification each time a new route is advertised to AWS from on premises over Direct Connect. What should a network engineer do to meet these requirements?

- A. Enable Amazon CloudWatch metrics on Direct Connect to track the received route
- B. Configure a CloudWatch alarm to send notifications when routes change.
- C. Onboard Transit Gateway Network Manager to Amazon CloudWatch Logs Insight
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to send notifications when routes change.
- E. Configure an AWS Lambda function to periodically check the routes on the Direct Connect gateway and to send notifications when routes change.
- F. Enable Amazon CloudWatch Logs on the transit VIFs to track the received route
- G. Create a metric filter Set an alarm on the filter to send notifications when routes change.

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/network-manager/latest/cloudwan/cloudwan-cloudwatch-events.html>

To receive notification each time a new route is advertised to AWS from on premises over Direct Connect, a network engineer should onboard Transit Gateway Network Manager to Amazon CloudWatch Logs Insights and use Amazon EventBridge (Amazon CloudWatch Events) to send notifications when routes change (Option B). This solution allows for real-time monitoring of route changes and automatic notification when new routes are advertised.

#### NEW QUESTION 98

A company is deploying a new application on AWS. The application uses dynamic multicasting. The company has five VPCs that are all attached to a transit gateway Amazon EC2 instances in each VPC need to be able to register dynamically to receive a multicast transmission. How should a network engineer configure the AWS resources to meet these requirements?

- A. Create a static source multicast domain within the transit gatewa
- B. Associate the VPCs and applicable subnets with the multicast domai
- C. Register the multicast senders' network interface with the multicast domai
- D. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
- E. Create a static source multicast domain within the transit gatewa
- F. Associate the VPCs and applicable subnets with the multicast domai
- G. Register the multicast senders' network interface with the multicast domai
- H. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.
- I. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domai
- J. Register the multicast senders' network interface with the multicast domai
- K. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
- L. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domai
- M. Register the multicast senders' network interface with the multicast domai
- N. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.

**Answer: C**

#### NEW QUESTION 102

A company's AWS architecture consists of several VPCs. The VPCs include a shared services VPC and several application VPCs. The company has established network connectivity from all VPCs to the on-premises DNS servers.

Applications that are deployed in the application VPCs must be able to resolve DNS for internally hosted domains on premises. The applications also must be able to resolve local VPC domain names and domains that are hosted in Amazon Route 53 private hosted zones.

What should a network engineer do to meet these requirements?

- A. Create a new Route 53 Resolver inbound endpoint in the shared services VP
- B. Create forwarding rules for the on-premises hosted domain
- C. Associate the rules with the new Resolver endpoint and each application VP
- D. Update each application VPC's DHCP configuration to point DNS resolution to the new Resolver endpoint.
- E. Create a new Route 53 Resolver outbound endpoint in the shared services VP
- F. Create forwarding rules for the on-premises hosted domain
- G. Associate the rules with the new Resolver endpoint and each application VPC.
- H. Create a new Route 53 Resolver outbound endpoint in the shared services VPCreate forwarding rules for the on-premises hosted domain
- I. Associate the rules with the new Resolver endpoint and each application VPUpdate each application VPC's DHCP configuration to point DNS resolution to the new Resolver endpoint.
- J. Create a new Route 53 Resolver inbound endpoint in the shared services VP
- K. Create forwarding rules for the on-premises hosted domain
- L. Associate the rules with the new Resolver endpoint and each application VPC.

**Answer: B**

**Explanation:**

Creating a new Route 53 Resolver outbound endpoint in the shared services VPC would enable forwarding of DNS queries from the VPC to on-premises1. Creating forwarding rules for the on-premises hosted domains would enable specifying which domain names are forwarded to the on-premises DNS servers2. Associating the rules with the new Resolver endpoint and each application VPC would enable applying the rules to the VPCs2. This solution would not affect the default DNS resolution behavior of Route 53 Resolver for local VPC domain names and domains that are hosted in Route 53 private hosted zones3.

**NEW QUESTION 104**

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

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

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

- A. Create a self-signed certificate for service.example.co
- B. Import the certificate into AWS Certificate Manager (ACM). Configure CloudFront to use this imported SSL/TLS certificat
- C. Change the default behavior to redirect HTTP to HTTPS.
- D. Create a certificate for service.example.com by using AWS Certificate Manager (ACM). Configure CloudFront to use this custom SSL/TLS certificat
- E. Change the default behavior to redirect HTTP to HTTPS.
- F. Create a certificate with any domain name by using AWS Certificate Manager (ACM) for the EC2 instance
- G. Configure the backend to use this certificate for its HTTPS listene
- H. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its target
- I. Attach the existing Auto Scaling group to this new target group.
- J. Create a public certificate from a third-party certificate provider with any domain name for the EC2 instance
- K. Configure the backend to use this certificate for its HTTPS listene
- L. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its target
- M. Attach the existing Auto Scaling group to this new target group.
- N. Create a certificate for service-alb.example.com by using AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the service-alb.example.com ACM certificat
- O. Modify the CloudFront origin to use the HTTPS protocol onl
- P. Delete the HTTP listener on the ALB.
- Q. Create a self-signed certificate for service-alb.example.co
- R. Import the certificate into AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the imported service-alb.example.com ACM certificat
- S. Modify the CloudFront origin to use the HTTPS protocol onl
- T. Delete the HTTP listener on the ALB.

**Answer:** BDE

**NEW QUESTION 108**

A company is migrating an application from on premises to AWS. The company will host the application on Amazon EC2 instances that are deployed in a single VPC. During the migration period, DNS queries from the EC2 instances must be able to resolve names of on-premises servers. The migration is expected to take 3 months After the 3-month migration period, the resolution of on-premises servers will no longer be needed.

What should a network engineer do to meet these requirements with the LEAST amount of configuration?

- A. Set up an AWS Site-to-Site VPN connection between on premises and AW
- B. Deploy an Amazon Route 53 Resolver outbound endpoint in the Region that is hosting the VPC.
- C. Set up an AWS Direct Connect connection with a private VI
- D. Deploy an Amazon Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint in the Region that is hosting the VPC.
- E. Set up an AWS Client VPN connection between on premises and AW
- F. Deploy an Amazon Route 53 Resolver inbound endpoint in the VPC.
- G. Set up an AWS Direct Connect connection with a public VI
- H. Deploy an Amazon Route 53 Resolver inbound endpoint in the Region that is hosting the VP
- I. Use the IP address that is assigned to the endpoint for connectivity to the on-premises DNS servers.

**Answer:** A

**Explanation:**

Setting up an AWS Site-to-Site VPN connection between on premises and AWS would enable a secure and encrypted connection over the public internet1.

Deploying an Amazon Route 53 Resolver outbound endpoint in the Region that is hosting the VPC would enable forwarding of DNS queries for on-premises servers to the on-premises DNS servers2. This would allow EC2 instances in the VPC to resolve names of on-premises servers during the migration period. After the migration period, the Route 53 Resolver outbound endpoint can be deleted with minimal configuration changes.

**NEW QUESTION 110**

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

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

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

- A. Create an Amazon DynamoDB table to maintain all IP address ranges and security groups that need to be update
- B. Update the DynamoDB table with the new IP address range when the company adds a new partne
- C. Invoke an AWS Lambda function to read new IP address ranges and security groups from the DynamoDB table to update the security group
- D. Deploy this solution in all accounts.
- E. Create a new prefix lis
- F. Add all allowed IP address ranges to the prefix lis
- G. Use Amazon EventBridge (Amazon CloudWatch Events) rules to invoke an AWS Lambda function to update security groups whenever a new IP address range is added to the prefix lis



- H. Deploy this solution in all accounts.
- I. Create a new prefix list
- J. Add all allowed IP address ranges to the prefix list
- K. Share the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM). Update security groups to use the prefix list instead of the partner IP address range
- L. Update the prefix list with the new IP address range when the company adds a new partner.
- M. Create an Amazon S3 bucket to maintain all IP address ranges and security groups that need to be updated
- N. Update the S3 bucket with the new IP address range when the company adds a new partner
- O. Invoke an AWS Lambda function to read new IP address ranges and security groups from the S3 bucket to update the security group
- P. Deploy this solution in all accounts.

**Answer: C**

**Explanation:**

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

**NEW QUESTION 115**

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

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

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

Which solution will meet these requirements?

- A. Create VPN attachments between the two transit gateways
- B. Configure the VPN attachments to use BGP routing between the two transit gateways.
- C. Peer the transit gateways in each Region
- D. Configure routing between the two transit gateways for each Region's IP addresses.
- E. Create a VPN server in a VPC in each Region
- F. Update the routing to point to the VPN servers for the IP addresses in alternate Regions.
- G. Attach the VPCs in us-west-1 to the transit gateway in ap-southeast-2.

**Answer: B**

**Explanation:**

Peering the transit gateways in each region would establish a private network connection between the two regions, allowing the company to route traffic between the VPCs in different regions without going over the public internet. This would help minimize latency and maximize bandwidth while reducing the operational overhead of managing multiple VPN connections.

**NEW QUESTION 118**

A network engineer has deployed an Amazon EC2 instance in a private subnet in a VPC. The VPC has no public subnet. The EC2 instance hosts application code that sends messages to an Amazon Simple Queue Service (Amazon SQS) queue. The subnet has the default network ACL with no modification applied. The EC2 instance has the default security group with no modification applied.

The SQS queue is not receiving messages.

Which of the following are possible causes of this problem? (Choose two.)

- A. The EC2 instance is not attached to an IAM role that allows write operations to Amazon SQS.
- B. The security group is blocking traffic to the IP address range used by Amazon SQS
- C. There is no interface VPC endpoint configured for Amazon SQS
- D. The network ACL is blocking return traffic from Amazon SQS
- E. There is no route configured in the subnet route table for the IP address range used by Amazon SQS

**Answer: CE**

**NEW QUESTION 121**

A network engineer is designing the architecture for a healthcare company's workload that is moving to the AWS Cloud. All data to and from the on-premises environment must be encrypted in transit. All traffic also must be inspected in the cloud before the traffic is allowed to leave the cloud and travel to the on-premises environment or to the internet.

The company will expose components of the workload to the internet so that patients can reserve appointments. The architecture must secure these components and protect them against DDoS attacks. The architecture also must provide protection against financial liability for services that scale out during a DDoS event.

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

- A. Use Traffic Mirroring to copy all traffic to a fleet of traffic capture appliances.
- B. Set up AWS WAF on all network components.
- C. Configure an AWS Lambda function to create Deny rules in security groups to block malicious IP addresses.
- D. Use AWS Direct Connect with MACsec support for connectivity to the cloud.
- E. Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection.
- F. Configure AWS Shield Advanced and ensure that it is configured on all public assets.

**Answer: DEF**

**Explanation:**

To meet the requirements for the healthcare company's workload that is moving to the AWS Cloud, the network engineer should take the following steps:

- Use AWS Direct Connect with MACsec support for connectivity to the cloud to ensure that all data to and from the on-premises environment is encrypted in transit (Option D).
  - Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection to inspect all traffic in the cloud before it is allowed to leave (Option E).
  - Configure AWS Shield Advanced and ensure that it is configured on all public assets to secure components exposed to the internet against DDoS attacks and provide protection against financial liability for services that scale out during a DDoS event (Option F).
- These steps will help ensure that all data is encrypted in transit, all traffic is inspected before leaving the cloud, and components exposed to the internet are secured against DDoS attacks.

#### NEW QUESTION 122

A company is migrating an existing application to a new AWS account. The company will deploy the application in a single AWS Region by using one VPC and multiple Availability Zones. The application will run on Amazon EC2 instances. Each Availability Zone will have several EC2 instances. The EC2 instances will be deployed in private subnets.

The company's clients will connect to the application by using a web browser with the HTTPS protocol. Inbound connections must be distributed across the Availability Zones and EC2 instances. All connections from the same client session must be connected to the same EC2 instance. The company must provide end-to-end encryption for all connections between the clients and the application by using the application SSL certificate.

Which solution will meet these requirements?

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

**Answer:** A

#### NEW QUESTION 124

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### ANS-C01 Practice Exam Features:

- \* ANS-C01 Questions and Answers Updated Frequently
- \* ANS-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* ANS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* ANS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The ANS-C01 Practice Test Here](#)**