

# Amazon-Web-Services

## Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam



### NEW QUESTION 1

A company is using custom DNS servers that run BIND for name resolution in its VPCs. The VPCs are deployed across multiple AWS accounts that are part of the same organization in AWS Organizations. All the VPCs are connected to a transit gateway. The BIND servers are running in a central VPC and are configured to forward all queries for an on-premises DNS domain to DNS servers that are hosted in an on-premises data center. To ensure that all the VPCs use the custom DNS servers, a network engineer has configured a VPC DHCP options set in all the VPCs that specifies the custom DNS servers to be used as domain name servers.

Multiple development teams in the company want to use Amazon Elastic File System (Amazon EFS). A development team has created a new EFS file system but cannot mount the file system to one of its Amazon EC2 instances. The network engineer discovers that the EC2 instance cannot resolve the IP address for the EFS mount point fs-33444567d.efs.us-east-1.amazonaws.com. The network engineer needs to implement a solution so that development teams throughout the organization can mount EFS file systems.

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

- A. Configure the BIND DNS servers in the central VPC to forward queries forefs.us-east-1.amazonaws.com to the Amazon provided DNS server (169.254.169.253).
- B. Create an Amazon Route 53 Resolver outbound endpoint in the central VP
- C. Update all the VPC DHCP options sets to use AmazonProvidedDNS for name resolution.
- D. Create an Amazon Route 53 Resolver inbound endpoint in the central VPCUpdate all the VPC DHCP options sets to use the Route 53 Resolver inbound endpoint in the central VPC for name resolution.
- E. Create an Amazon Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS server
- F. Share the rule with the organization by using AWS Resource Access Manager (AWS RAM). Associate the rule with all the VPCs.
- G. Create an Amazon Route 53 private hosted zone for the efs.us-east-1.amazonaws.com domain. Associate the private hosted zone with the VPC where the EC2 instance is deploye
- H. Create an A record for fs-33444567d.efs.us-east-1.amazonaws.com in the private hosted zon
- I. Configure the A record to return the mount target of the EFS mount point.

**Answer: BD**

#### Explanation:

Option B suggests using Amazon Route 53 Resolver outbound endpoint, which would replace the existing BIND DNS servers with the AmazonProvidedDNS for name resolution. However, the scenario specifically mentions that the company is using custom DNS servers that run BIND for name resolution in its VPCs, so this solution would not work. Option D suggests creating a Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS servers, which would not address the issue of resolving the EFS mount point. The problem is not with resolving queries for the on-premises domain, but rather with resolving the IP address for the EFS mount point.

### NEW QUESTION 2

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 NLUUse 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 AL
- G. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator
- H. Place the EC2 instances behind an Amazon CloudFront distributio
- I. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

**Answer: B**

### NEW QUESTION 3

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?

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

### NEW QUESTION 4

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 this centralized 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 5

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

A network engineer must implement connectivity between VPC-B and the on-premises data center in Dublin. Which solutions will meet these requirements? (Choose two.)

- A. Configure inter-Region VPC peering between VPC-A and VPC-
- B. Add the required VPC peering route
- C. Add the VPC-B CIDR block in the allowed prefixes on the Direct Connect gateway association.
- D. Associate TGW-B with the Direct Connect gateway
- E. Advertise the VPC-B CIDR block under the allowed prefixes.
- F. Configure another transit VIF on the Direct Connect connection and associate TGW-
- G. Advertise the VPC-B CIDR block under the allowed prefixes.
- H. Configure inter-Region transit gateway peering between TGW-A and TGW-
- I. Add the peering routes in the transit gateway route table
- J. Add both the VPC-A and the VPC-B CIDR block under the allowed prefix list in the Direct Connect gateway association.
- K. Configure an AWS Site-to-Site VPN connection over the transit VIF to TGW-B as a VPN attachment.

**Answer:** BC

#### Explanation:

\* B. Associate TGW-B with the Direct Connect gateway. Advertise the VPC-B CIDR block under the allowed prefixes. This will allow traffic from VPC-B to be sent over the Direct Connect connection to the on-premises data center via TGW-B. C. Configure another transit VIF on the Direct Connect connection and associate TGW-B. Advertise the VPC-B CIDR block under the allowed prefixes. This will enable the use of the Direct Connect connection for VPC-B's traffic by connecting TGW-B to the Direct Connect gateway.

#### NEW QUESTION 6

A company has deployed a web application on AWS. The web application uses an Application Load Balancer (ALB) across multiple Availability Zones. The targets of the ALB are AWS Lambda functions. The web application also uses Amazon CloudWatch metrics for monitoring.

Users report that parts of the web application are not loading properly. A network engineer needs to troubleshoot the problem. The network engineer enables access logging for the ALB.

What should the network engineer do next to determine which errors the ALB is receiving?

- A. Send the logs to Amazon CloudWatch Log
- B. Review the ALB logs in CloudWatch Insights to determine which error messages the ALB is receiving.
- C. Configure the Amazon S3 bucket destination
- D. Use Amazon Athena to determine which error messages the ALB is receiving.
- E. Configure the Amazon S3 bucket destination
- F. After Amazon CloudWatch Logs pulls the ALB logs from the S3 bucket automatically, review the logs in CloudWatch Logs to determine which error messages the ALB is receiving.
- G. Send the logs to Amazon CloudWatch Log
- H. Use the Amazon Athena CloudWatch Connector to determine which error messages the ALB is receiving.

**Answer:** B

#### Explanation:

Access logs is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logs for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify as compressed files. You can disable access logs at any time. <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

#### NEW QUESTION 7

A real estate company is building an internal application so that real estate agents can upload photos and videos of various properties. The application will store these photos and videos in an Amazon S3 bucket as objects and will use Amazon DynamoDB to store corresponding metadata. The S3 bucket will be configured

to publish all PUT events for new object uploads to an Amazon Simple Queue Service (Amazon SQS) queue.

A compute cluster of Amazon EC2 instances will poll the SQS queue to find out about newly uploaded objects. The cluster will retrieve new objects, perform proprietary image and video recognition and classification update metadata in DynamoDB and replace the objects with new watermarked objects. The company does not want public IP addresses on the EC2 instances.

Which networking design solution will meet these requirements MOST cost-effectively as application usage increases?

- A. Place the EC2 instances in a public subne
- B. Disable the Auto-assign Public IP option while launching the EC2 instance
- C. Create an internet gatewa
- D. Attach the internet gateway to the VP
- E. In the public subnet's route table, add a default route that points to the internet gateway.
- F. Place the EC2 instances in a private subne
- G. Create a NAT gateway in a public subnet in the same Availability Zon
- H. Create an internet gatewa
- I. Attach the internet gateway to the VP
- J. In the public subnet's route table, add a default route that points to the internet gateway
- K. Place the EC2 instances in a private subne
- L. Create an interface VPC endpoint for Amazon SQ
- M. Create gateway VPC endpoints for Amazon S3 and DynamoDB.
- N. Place the EC2 instances in a private subne
- O. Create a gateway VPC endpoint for Amazon SQS.Create interface VPC endpoints for Amazon S3 and DynamoDB.

**Answer: C**

### NEW QUESTION 8

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 9

A company is deploying third-party firewall appliances for traffic inspection and NAT capabilities in its VPC. The VPC is configured with private subnets and public subnets. The company needs to deploy the firewall appliances behind a load balancer.

Which architecture will meet these requirements MOST cost-effectively?

- A. Deploy a Gateway Load Balancer with the firewall appliances as target
- B. Configure the firewall appliances with a single network interface in a private subne
- C. Use a NAT gateway to send the traffic to the internet after inspection.
- D. Deploy a Gateway Load Balancer with the firewall appliances as target
- E. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subne
- F. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.
- G. Deploy a Network Load Balancer with the firewall appliances as target
- H. Configure the firewall appliances with a single network interface in a private subne
- I. Use a NAT gateway to send the traffic to the internet after inspection.
- J. Deploy a Network Load Balancer with the firewall appliances as target
- K. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subne
- L. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.

**Answer: B**

### NEW QUESTION 10

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 subnets
- C. For the second network interface, select one of the public subnets
- 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 10**

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 15**

Your company runs an application for the US market in the us-east-1 AWS region. This application uses proprietary TCP and UDP protocols on Amazon Elastic Compute Cloud (EC2) instances. End users run a

real-time, front-end application on their local PCs. This front-end application knows the DNS hostname of the service.

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

How should you use AWS services to meet these requirements?

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

**Answer: B**

**NEW QUESTION 18**

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 19**

A network engineer needs to update a company's hybrid network to support IPv6 for the upcoming release of a new application. The application is hosted in a VPC in the AWS Cloud. The company's current AWS infrastructure includes VPCs that are connected by a transit gateway. The transit gateway is connected to the on-premises network by AWS Direct Connect and AWS Site-to-Site VPN. The company's on-premises devices have been updated to support the new IPv6 requirements.

The company has enabled IPv6 for the existing VPC by assigning a new IPv6 CIDR block to the VPC and by assigning IPv6 to the subnets for dual-stack support. The company has launched new Amazon EC2 instances for the new application in the updated subnets.

When updating the hybrid network to support IPv6 the network engineer must avoid making any changes to the current infrastructure. The network engineer also must block direct access to the instances' new IPv6 addresses from the internet. However, the network engineer must allow outbound internet access from the instances.

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

- A. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address
- B. Create a new VPN connection that supports IPv6 connectivity
- C. Add an egress-only internet gateway
- D. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices
- E. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address
- F. Update the existing VPN connection to support IPv6 connectivity
- G. Add an egress-only internet gateway
- H. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
- I. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address
- J. Create a new VPN connection that supports IPv6 connectivity
- K. Add an egress-only internet gateway
- L. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
- M. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address
- N. Create a new VPN connection that supports IPv6 connectivity
- O. Add a NAT gateway
- P. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.

**Answer: B**

#### NEW QUESTION 21

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 26

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 VPC
- 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 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 VPC
- 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. Creating a link aggregation group (LAG) with the existing and new Direct Connect connections would provide resiliency and redundancy for the AWS connectivity.

#### NEW QUESTION 29

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 31

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 34

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:

AWSTemplateFormatVersion: 2010-09-09 Parameters:

OriginatingVPCId: 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::EC2::VPCPeeringConnection'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 37

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 VP
- 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 VP
- 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 41**

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 44**

A company is hosting an application on Amazon EC2 instances behind an Application Load Balancer. The instances are in an Amazon EC2 Auto Scaling group. Because of a recent change to a security group, external users cannot access the application. A network engineer needs to prevent this downtime from happening again. The network engineer must implement a solution that remediates noncompliant changes to security groups. Which solution will meet these requirements?

- A. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuration
- B. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.
- C. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration
- D. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.
- E. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuration
- F. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.
- G. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration
- H. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.

**Answer: D**

**Explanation:**

Configuring an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration would enable evaluation of the compliance status of the security groups based on predefined or custom rules<sup>3</sup>. Creating an AWS Systems Manager Automation runbook to remediate noncompliant security groups would enable automation of the remediation process<sup>2</sup>. Additionally, configuring AWS Config to trigger the runbook when a noncompliant change is detected would enable timely and consistent remediation of security group changes.

**NEW QUESTION 49**

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 bucket
- 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 52

A network engineer must provide additional safeguards to protect encrypted data at Application Load Balancers (ALBs) through the use of a unique random session key.

What should the network engineer do to meet this requirement?

- A. Change the ALB security policy to a policy that supports TLS 1.2 protocol only
- B. Use AWS Key Management Service (AWS KMS) to encrypt session keys
- C. Associate an AWS WAF web ACL with the ALB
- D. and create a security rule to enforce forward secrecy (FS)
- E. Change the ALB security policy to a policy that supports forward secrecy (FS)

**Answer: D**

#### NEW QUESTION 54

A company is planning to create a service that requires encryption in transit. The traffic must not be decrypted between the client and the backend of the service. The company will implement the service by using the gRPC protocol over TCP port 443. The service will scale up to thousands of simultaneous connections. The backend of the service will be hosted on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with the Kubernetes Cluster Autoscaler and the Horizontal Pod Autoscaler configured. The company needs to use mutual TLS for two-way authentication between the client and the backend.

Which solution will meet these requirements?

- A. Install the AWS Load Balancer Controller for Kubernetes
- B. Using that controller, configure a Network Load Balancer with a TCP listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
- C. Install the AWS Load Balancer Controller for Kubernetes
- D. Using that controller, configure an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
- E. Create a target group
- F. Add the EKS managed node group's Auto Scaling group as a target. Create an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the target group.
- G. Create a target group
- H. Add the EKS managed node group's Auto Scaling group as a target
- I. Create a Network Load Balancer with a TLS listener on port 443 to forward traffic to the target group.

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html#target-gro>

#### NEW QUESTION 55

A company has deployed its AWS environment in a single AWS Region. The environment consists of a few hundred application VPCs, a shared services VPC, and a VPN connection to the company's on-premises environment. A network engineer needs to implement a transit gateway with the following requirements:

- Application VPCs must be isolated from each other.
- Bidirectional communication must be allowed between the application VPCs and the on-premises network.
- Bidirectional communication must be allowed between the application VPCs and the shared services VPC. The network engineer creates the transit gateway with options disabled for default route table association and default route table propagation. The network engineer also creates the VPN attachment for the on-premises network and creates the VPC attachments for the application VPCs and the shared services VPC.

The network engineer must meet all the requirements for the transit gateway by designing a solution that needs the least number of transit gateway route tables.

Which combination of actions should the network engineer perform to accomplish this goal?(Choose two.)

- A. Configure a separate transit gateway route table for on premise
- B. Associate the VPN attachment with this transit gateway route table
- C. Propagate all application VPC attachments to this transit gateway route table.
- D. Configure a separate transit gateway route table for each application VPC
- E. Associate each application VPC attachment with its respective transit gateway route table
- F. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.
- G. Configure a separate transit gateway route table for all application VPC
- H. Associate all application VPCs with this transit gateway route table
- I. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.
- J. Configure a separate transit gateway route table for the shared services VPC
- K. Associate the shared services VPC attachment with this transit gateway route table
- L. Propagate all application VPC attachments to this transit gateway route table.
- M. Configure a separate transit gateway route table for on premises and the shared services VPC
- N. Associate the VPN attachment and the shared services VPC attachment with this transit gateway route table
- O. Propagate all application VPC attachments to this transit gateway route table.

**Answer: BD**

#### NEW QUESTION 57

An organization is using a VPC endpoint for Amazon S3. When the security group rules for a set of instances were initially configured, access was restricted to allow traffic only to the IP addresses of the Amazon S3 API endpoints in the region from the published JSON file. The application was working properly, but now is logging a growing number of timeouts when connecting with Amazon S3. No internet gateway is configured for the VPC.

Which solution will fix the connectivity failures with the LEAST amount of effort?

- A. Create a Lambda function to update the security group based on AmazonIPSpaceChanged notifications.
- B. Update the VPC routing to direct Amazon S3 prefix-list traffic to the VPC endpoint using the route table APIs.
- C. Update the application server's outbound security group to use the prefix-list for Amazon S3 in the same region.
- D. Create an additional VPC endpoint for Amazon S3 in the same route table to scale the concurrent connections to Amazon.

**Answer: C**

**Explanation:**

<https://aws.amazon.com/blogs/aws/subscribe-to-aws-public-ip-address-changes-via-amazon-sns/>

**NEW QUESTION 59**

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 the existing subnets.

**Answer: B**

**NEW QUESTION 62**

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 64**

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 appliance
- B. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway
- C. Configure BGP over the IPsec VPN connections
- D. Assign a new CIDR block to the transit gateway
- E. Create a new VPC for the SD-WAN hub virtual appliance
- F. Attach the new VPC to the transit gateway with a VPC attachment
- G. Add a transit gateway Connect attachment
- 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 appliance
- K. Attach the new VPC to the transit gateway with a VPC attachment
- L. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway
- M. Configure BGP over the IPsec VPN connections.
- N. Assign a new CIDR block to the transit gateway
- O. Create a new VPC for the SD-WAN hub virtual appliance
- P. Attach the new VPC to the transit gateway with a VPC attachment
- Q. Add a transit gateway Connect attachment
- 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 65**

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 IP 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 69**

A company's network engineer needs to design a new solution to help troubleshoot and detect network anomalies. The network engineer has configured Traffic Mirroring. However, the mirrored traffic is overwhelming the Amazon EC2 instance that is the traffic mirror target. The EC2 instance hosts tools that the company's security team uses to analyze the traffic. The network engineer needs to design a highly available solution that can scale to meet the demand of the mirrored traffic.

Which solution will meet these requirements?

- A. Deploy a Network Load Balancer (NLB) as the traffic mirror target
- B. Behind the NLB
- C. Deploy a fleet of EC2 instances in an Auto Scaling group
- D. Use Traffic Mirroring as necessary.
- E. Deploy an Application Load Balancer (ALB) as the traffic mirror target
- F. Behind the ALB, deploy a fleet of EC2 instances in an Auto Scaling group
- G. Use Traffic Mirroring only during non-business hours.
- H. Deploy a Gateway Load Balancer (GLB) as the traffic mirror target
- I. Behind the GLB
- J. Deploy a fleet of EC2 instances in an Auto Scaling group
- K. Use Traffic Mirroring as necessary.
- L. Deploy an Application Load Balancer (ALB) with an HTTPS listener as the traffic mirror target
- M. Behind the ALB
- N. Deploy a fleet of EC2 instances in an Auto Scaling group
- O. Use Traffic Mirroring only during active events or business hours.

**Answer: A**

**NEW QUESTION 72**

A company's network engineer builds and tests network designs for VPCs in a development account. The company needs to monitor the changes that are made to network resources and must ensure strict compliance with network security policies. The company also needs access to the historical configurations of network resources.

Which solution will meet these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a custom pattern to monitor the account for changes
- B. Configure the rule to invoke an AWS Lambda function to identify noncompliant resources
- C. Update an Amazon DynamoDB table with the changes that are identified.
- D. Create custom metrics from Amazon CloudWatch logs
- E. Use the metrics to invoke an AWS Lambda function to identify noncompliant resources
- F. Update an Amazon DynamoDB table with the changes that are identified.
- G. Record the current state of network resources by using AWS Config
- H. Create rules that reflect the desired configuration settings
- I. Set remediation for noncompliant resources.
- J. Record the current state of network resources by using AWS Systems Manager Inventory
- K. Use Systems Manager State Manager to enforce the desired configuration settings and to carry out remediation for noncompliant resources.

**Answer: C**

**Explanation:**

Recording the current state of network resources by using AWS Config would enable auditing and assessment of resource configurations and compliance.  
Creating rules that reflect the desired configuration settings would enable evaluation of whether the network resources comply with network security policies.  
Setting remediation for noncompliant resources would enable automatic correction of undesired configurations.

**NEW QUESTION 73**

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 78

A company deploys a new web application on Amazon EC2 instances. The application runs in private subnets in three Availability Zones behind an Application Load Balancer (ALB). Security auditors require encryption of all connections. The company uses Amazon Route 53 for DNS and uses AWS Certificate Manager (ACM) to automate SSL/TLS certificate provisioning. SSL/TLS connections are terminated on the ALB.

The company tests the application with a single EC2 instance and does not observe any problems. However, after production deployment, users report that they can log in but that they cannot use the application. Every new web request restarts the login process.

What should a network engineer do to resolve this issue?

- A. Modify the ALB listener configuration.
- B. Edit the rule that forwards traffic to the target group.
- C. Change the rule to enable group-level stickiness.
- D. Set the duration to the maximum application session length.
- E. Replace the ALB with a Network Load Balancer.
- F. Create a TLS listener.
- G. Create a new target group with the protocol type set to TLS. Register the EC2 instance.
- H. Modify the target group configuration by enabling the stickiness attribute.
- I. Modify the ALB target group configuration by enabling the stickiness attribute.
- J. Use an application-based cookie.
- K. Set the duration to the maximum application session length.
- L. Remove the ALB.
- M. Create an Amazon Route 53 rule with a failover routing policy for the application name.
- N. Configure ACM to issue certificates for each EC2 instance.

**Answer:** C

#### NEW QUESTION 80

A company is developing an application in which IoT devices will report measurements to the AWS Cloud. The application will have millions of end users. The company observes that the IoT devices cannot support DNS resolution. The company needs to implement an Amazon EC2 Auto Scaling solution so that the IoT devices can connect to an application endpoint without using DNS.

Which solution will meet these requirements MOST cost-effectively?

- A. Use an Application Load Balancer (ALB)-type target group for a Network Load Balancer (NLB). Create an EC2 Auto Scaling group.
- B. Attach the Auto Scaling group to the ALB.
- C. Set up the IoT devices to connect to the IP addresses of the NLB.
- D. Use an AWS Global Accelerator accelerator with an Application Load Balancer (ALB) endpoint.
- E. Create an EC2 Auto Scaling group.
- F. Attach the Auto Scaling group to the ALB. Set up the IoT devices to connect to the IP addresses of the accelerator.
- G. Use a Network Load Balancer (NLB). Create an EC2 Auto Scaling group.
- H. Attach the Auto Scaling group to the NLB.
- I. Set up the IoT devices to connect to the IP addresses of the NLB.
- J. Use an AWS Global Accelerator accelerator with a Network Load Balancer (NLB) endpoint.
- K. Create an EC2 Auto Scaling group.
- L. Attach the Auto Scaling group to the NLB.
- M. Set up the IoT devices to connect to the IP addresses of the accelerator.

**Answer:** D

#### Explanation:

AWS Global Accelerator can provide static IP addresses that the IoT devices can connect to without using DNS. It can also route traffic over the AWS global network and improve performance and availability for the IoT devices. An NLB can provide end-to-end encryption for HTTPS traffic by using TLS as a target group protocol and terminating SSL connections at the load balancer level. An NLB can also support session affinity (sticky sessions) with TCP connections.

#### NEW QUESTION 85

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 90**

A company has multiple AWS accounts. Each account contains one or more VPCs. A new security guideline requires the inspection of all traffic between VPCs. The company has deployed a transit gateway that provides connectivity between all VPCs. The company also has deployed a shared services VPC with Amazon EC2 instances that include IDS services for stateful inspection. The EC2 instances are deployed across three Availability Zones. The company has set up VPC associations and routing on the transit gateway. The company has migrated a few test VPCs to the new solution for traffic inspection. Soon after the configuration of routing, the company receives reports of intermittent connections for traffic that crosses Availability Zones. What should a network engineer do to resolve this issue?

- A. Modify the transit gateway VPC attachment on the shared services VPC by enabling cross-Availability Zone load balancing.
- B. Modify the transit gateway VPC attachment on the shared services VPC by enabling appliance mode support.
- C. Modify the transit gateway by selecting VPN equal-cost multi-path (ECMP) routing support.
- D. Modify the transit gateway by selecting multicast support.

**Answer:** B

**Explanation:**

To resolve the issue of intermittent connections for traffic that crosses Availability Zones after configuring routing for traffic inspection between VPCs using a transit gateway and EC2 instances with IDS services in a shared services VPC, a network engineer should modify the transit gateway VPC attachment on the shared services VPC by enabling appliance mode support (Option B). This will ensure that traffic is routed to the same EC2 instance for stateful inspection and prevent intermittent connections.

**NEW QUESTION 92**

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 certificate
- 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 certificate
- 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 listener
- 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 listener
- 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 certificate
- O. Modify the CloudFront origin to use the HTTPS protocol only
- 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 certificate
- S. Modify the CloudFront origin to use the HTTPS protocol only
- T. Delete the HTTP listener on the ALB.

**Answer:** BDE

**NEW QUESTION 94**

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 AWS
- 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 VIF
- 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 AWS
- F. Deploy an Amazon Route 53 Resolver inbound endpoint in the VPC.
- G. Set up an AWS Direct Connect connection with a public VIF
- H. Deploy an Amazon Route 53 Resolver inbound endpoint in the Region that is hosting the VPC
- 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 internet<sup>1</sup>. 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 servers<sup>2</sup>. 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 96**

A company uses AWS Direct Connect to connect its corporate network to multiple VPCs in the same AWS account and the same AWS Region. Each VPC uses its own private VIF and its own virtual LAN on the Direct Connect connection. The company has grown and will soon surpass the limit of VPCs and private VIFs for each connection.

What is the MOST scalable way to add VPCs with on-premises connectivity?

- A. Provision a new Direct Connect connection to handle the additional VPC
- B. Use the new connection to connect additional VPCs.
- C. Create virtual private gateways for each VPC that is over the service quot
- D. Use AWS Site-to-Site VPNto connect the virtual private gateways to the corporate network.
- E. Create a Direct Connect gateway, and add virtual private gateway associations to the VPC
- F. Configure a private VIF to connect to the corporate network.
- G. Create a transit gateway, and attach the VPC
- H. Create a Direct Connect gateway, and associate it with the transit gatewa
- I. Create a transit VIF to the Direct Connect gateway.

**Answer:** D

**Explanation:**

When a company requires connectivity to multiple VPCs over AWS Direct Connect, a scalable solution is to use a transit gateway. A transit gateway is a hub that can interconnect multiple VPCs and VPN connections. The VPCs can communicate with each other over the transitgateway, and on-premises networks can communicate with the VPCs through the Direct Connect gateway. This solution provides a central point of management and simplifies the configuration of network routing. By associating the Direct Connect gateway with the transit gateway, traffic between the VPCs and the on-premises network can be routed through the Direct Connect connection.

**NEW QUESTION 99**

A company recently migrated its Amazon EC2 instances to VPC private subnets to satisfy a security compliance requirement. The EC2 instances now use a NAT gateway for internet access. After the migration, some long-running database queries from private EC2 instances to a publicly accessible third-party database no longer receive responses. The database query logs reveal that the queries successfully completed after 7 minutes but that the client EC2 instances never received the response.

Which configuration change should a network engineer implement to resolve this issue?

- A. Configure the NAT gateway timeout to allow connections for up to 600 seconds.
- B. Enable enhanced networking on the client EC2 instances.
- C. Enable TCP keepalive on the client EC2 instances with a value of less than 300 seconds.
- D. Close idle TCP connections through the NAT gateway.

**Answer:** C

**Explanation:**

When a TCP connection is idle for a long time, it may be terminated by network devices, including the NAT gateway. By enabling TCP keepalive, the client EC2 instances can periodically send packets to the third-party database to indicate that the connection is still active, preventing it from being terminated prematurely.

**NEW QUESTION 103**

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 108**

Your security team implements a host-based firewall on all of your Amazon Elastic Compute Cloud (EC2) instances to block all outgoing traffic. Exceptions must be requested for each specific requirement. Until you request a new rule, you cannot access the instance metadata service. Which firewall rule should you request to be added to your instances to allow instance metadata access?

- A. Inbound; Protocol tcp; Source [Instance's EIP]; Destination 169.254.169.254
- B. Inbound; Protocol tcp; Destination 169.254.169.254; Destination port 80
- C. Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 80
- D. Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 443

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html>  
To view all categories of instance metadata from within a running instance, use the following URI.  
<http://169.254.169.254/latest/meta-data/>

#### NEW QUESTION 109

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 114

.....

## **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 First Try
- \* ANS-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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