

Amazon-Web-Services

Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam



NEW QUESTION 1

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: A

NEW QUESTION 2

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 3

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 subnet
- B. Disable the Auto-assign Public IP option while launching the EC2 instance
- C. Create an internet gateway
- D. Attach the internet gateway to the VPC
- 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 subnet
- G. Create a NAT gateway in a public subnet in the same Availability Zone
- H. Create an internet gateway
- I. Attach the internet gateway to the VPC
- 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 subnet
- L. Create an interface VPC endpoint for Amazon SQS
- M. Create gateway VPC endpoints for Amazon S3 and DynamoDB.
- N. Place the EC2 instances in a private subnet
- O. Create a gateway VPC endpoint for Amazon SQS. Create interface VPC endpoints for Amazon S3 and DynamoDB.

Answer: C

NEW QUESTION 4

A company is using an AWS Site-to-Site VPN connection from the company's on-premises data center to a virtual private gateway in the AWS Cloud. Because of congestion, the company is experiencing availability and performance issues as traffic travels across the internet before the traffic reaches AWS. A network engineer must reduce these issues for the connection as quickly as possible with minimum administration effort.

Which solution will meet these requirements?

- A. Edit the existing Site-to-Site VPN connection by enabling acceleration
- B. Stop and start the VPN service on the customer gateway for the new setting to take effect.
- C. Configure a transit gateway in the same AWS Region as the existing virtual private gateway

- D. Create a new accelerated Site-to-Site VPN connectio
- E. Connect the new connection to the transit gateway by using a VPN attachmen
- F. Update the customer gateway device to use the new Site to Site VPN connectio
- G. Delete the existing Site-to-Site VPN connection
- H. Create a new accelerated Site-to-Site VPN connectio
- I. Connect the new Site-to-Site VPN connection to the existing virtual private gatewa
- J. Update the customer gateway device to use the new Site-to-Site VPN connectio
- K. Delete the existing Site-to-Site VPN connection.
- L. Create a new AWS Direct Connect connection with a private VIF between the on-premises data center and the AWS Clou
- M. Update the customer gateway device to use the new Direct Connect connectio
- N. Delete the existing Site-to-Site VPN connection.

Answer: B

NEW QUESTION 5

A data analytics company has a 100-node high performance computing (HPC) cluster. The HPC cluster is for parallel data processing and is hosted in a VPC in the AWS Cloud. As part of the data processing workflow, the HPC cluster needs to perform several DNS queries to resolve and connect to Amazon RDS databases, Amazon S3 buckets, and on-premises data stores that are accessible through AWS Direct Connect. The HPC cluster can increase in size by five to seven times during the company's peak event at the end of the year.

The company is using two Amazon EC2 instances as primary DNS servers for the VPC. The EC2 instances are configured to forward queries to the default VPC resolver for Amazon Route 53 hosted domains and to the on-premises DNS servers for other on-premises hosted domain names. The company notices job failures and finds that DNS queries from the HPC cluster nodes failed when the nodes tried to resolve RDS and S3 bucket endpoints.

Which architectural change should a network engineer implement to provide the DNS service in the MOST scalable way?

- A. Scale out the DNS service by adding two additional EC2 instances in the VP
- B. Reconfigure half of the HPC cluster nodes to use these new DNS server
- C. Plan to scale out by adding additional EC2instance-based DNS servers in the future as the HPC cluster size grows.
- D. Scale up the existing EC2 instances that the company is using as DNS server
- E. Change the instance size to the largest possible instance size to accommodate the current DNS load and theanticipated load in the future.
- F. Create Route 53 Resolver outbound endpoint
- G. Create Route 53 Resolver rules to forward queries to on-premises DNS servers for on premises hosted domain name
- H. Reconfigure the HPC cluster nodes to use the default VPC resolver instead of the EC2 instance-based DNS server
- I. Terminate the EC2 instances.
- J. Create Route 53 Resolver inbound endpoint
- K. Create rules on the on-premises DNS servers to forward queries to the default VPC resolve
- L. Reconfigure the HPC cluster nodes to forward all DNS queries to the on-premises DNS server
- M. Terminate the EC2 instances.

Answer: C

NEW QUESTION 6

A company is planning to use Amazon S3 to archive financial data. The data is currently stored in an on-premises data center. The company uses AWS Direct Connect with a Direct Connect gateway and a transit gateway to connect to the on-premises data center. The data cannot be transported over the public internet and must be encrypted in transit.

Which solution will meet these requirements?

- A. Create a Direct Connect public VI
- B. Set up an IPsec VPN connection over the public VIF to access Amazon S3. Use HTTPS for communication.
- C. Create an IPsec VPN connection over the transit VI
- D. Create a VPC and attach the VPC to the transit gatewa
- E. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.
- F. Create a VPC and attach the VPC to the transit gatewa
- G. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.
- H. Create a Direct Connect public VI
- I. Set up an IPsec VPN connection over the public VIF to the transit gatewa
- J. Create an attachment for Amazon S3. Use HTTPS for communication.

Answer: B

NEW QUESTION 7

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 the 0.02 * OUT, ergo VPC peering will be cheaper

NEW QUESTION 8

A company uses a 4 Gbps AWS Direct Connect dedicated connection with a link aggregation group (LAG) bundle to connect to five VPCs that are deployed in the

us-east-1 Region. Each VPC serves a different business unit and uses its own private VIF for connectivity to the on-premises environment. Users are reporting slowness when they access resources that are hosted on AWS.

A network engineer finds that there are sudden increases in throughput and that the Direct Connect connection becomes saturated at the same time for about an hour each business day. The company wants to know which business unit is causing the sudden increase in throughput. The network engineer must find out this information and implement a solution to resolve the problem.

Which solution will meet these requirements?

- A. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observe
- B. Create a new 10 Gbps dedicated connectio
- C. Shift traffic from the existing dedicated connection to the new dedicated connection.
- D. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observe
- E. Upgrade the bandwidth of the existing dedicated connection to 10 Gbps.
- F. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observe
- G. Upgrade the existing dedicated connection to a 5 Gbps hosted connection.
- H. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observe
- I. Create a new 10 Gbps dedicated connectio
- J. Shift traffic from the existing dedicated connection to the new dedicated connection.

Answer: A

Explanation:

To meet the requirements of finding out which business unit is causing the sudden increase in throughput and resolving the problem, the network engineer should review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observed (Option B). After identifying the VIF that is causing the issue, they can upgrade the bandwidth of the existing dedicated connection to 10 Gbps to resolve the problem (Option B).

NEW QUESTION 9

A global company operates all its non-production environments out of three AWS Regions: eu-west-1, us-east-1, and us-west-1. The company hosts all its production workloads in two on-premises data centers. The company has 60 AWS accounts and each account has two VPCs in each Region. Each VPC has a virtual private gateway where two VPN connections terminate for resilient connectivity to the data centers. The company has 360 VPN tunnels to each data center, resulting in high management overhead. The total VPN throughput for each Region is 500 Mbps.

The company wants to migrate the production environments to AWS. The company needs a solution that will simplify the network architecture and allow for future growth. The production environments will generate an additional 2 Gbps of traffic per Region back to the data centers. This traffic will increase over time. Which solution will meet these requirements?

- A. Set up an AWS Direct Connect connection from each data center to AWS in each Regio
- B. Create and attach private VIFs to a single Direct Connect gatewa
- C. Attach the Direct Connect gateway to all the VPC
- D. Remove the existing VPN connections that are attached directly to the virtual private gateways.
- E. Create a single transit gateway with VPN connections from each data cente
- F. Share the transit gateway with each account by using AWS Resource Access Manager (AWS RAM). Attach the transit gateway to each VP
- G. Remove the existing VPN connections that are attached directly to the virtual private gateways.
- H. Create a transit gateway in each Region with multiple newly commissioned VPN connections from each data cente
- I. Share the transit gateways with each account by using AWS Resource Access Manager (AWS RAM). In each Region, attach the transit gateway to each VPC
- J. Peer all the VPCs in each Region to a new VPC in each Region that will function as a centralized transit VP
- K. Create new VPN connections from each data center to the transit VPC
- L. Terminate the original VPN connections that are attached to all the original VPC
- M. Retain the new VPN connection to the new transit VPC in each Region.

Answer: C

NEW QUESTION 10

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 10

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 configuratio
- B. Edit the rule that forwards traffic to the target grou
- C. Change the rule to enable group-level stickines
- D. Set the duration to the maximum application session length.
- E. Replace the ALB with a Network Load Balance
- F. Create a TLS listene
- 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 attribut
- J. Use an application-based cooki
- K. Set the duration to the maximum application session length.
- L. Remove the AL
- M. Create an Amazon Route 53 rule with a failover routing policy for the application nam
- N. Configure ACM to issue certificates for each EC2 instance.

Answer: C

NEW QUESTION 13

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 gatewa
- B. Create a VPC attachment to each application VP
- 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 Marketplac
- G. Create a VPN attachment from each VPC to the transit VP
- 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 14

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 15

A company has two on-premises data center locations. There is a company-managed router at each data center. Each data center has a dedicated AWS Direct Connect connection to a Direct Connect gateway through a private virtual interface. The router for the first location is advertising 110 routes to the Direct Connect gateway by using BGP, and the router for the second location is advertising 60 routes to the Direct Connect gateway by using BGP. The Direct Connect gateway is attached to a company VPC through a virtual private gateway.

A network engineer receives reports that resources in the VPC are not reachable from various locations in either data center. The network engineer checks the VPC route table and sees that the routes from the first data center location are not being populated into the route table. The network engineer must resolve this issue in the most operationally efficient manner.

What should the network engineer do to meet these requirements?

- A. Remove the Direct Connect gateway, and create a new private virtual interface from each company router to the virtual private gateway of the VPC.
- B. Change the router configurations to summarize the advertised routes.
- C. Open a support ticket to increase the quota on advertised routes to the VPC route table.
- D. Create an AWS Transit Gatewa
- E. Attach the transit gateway to the VPC, and connect the Direct Connect gateway to the transit gateway.

Answer: B

Explanation:

"If you advertise more than 100 routes each for IPv4 and IPv6 over the BGP session, the BGP session will go into an idle state with the BGP session

DOWN."https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html

NEW QUESTION 16

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 19

An ecommerce company is hosting a web application on Amazon EC2 instances to handle continuously changing customer demand. The EC2 instances are part of an Auto Scaling group. The company wants to implement a solution to distribute traffic from customers to the EC2 instances. The company must encrypt all traffic at all stages between the customers and the application servers. No decryption at intermediate points is allowed.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB). Add an HTTPS listener to the ALB
- B. Configure the Auto Scaling group to register instances with the ALB's target group.
- C. Create an Amazon CloudFront distribution
- D. Configure the distribution with a custom SSL/TLS certificate
- E. Set the Auto Scaling group as the distribution's origin.
- F. Create a Network Load Balancer (NLB). Add a TCP listener to the NLB
- G. Configure the Auto Scaling group to register instances with the NLB's target group.
- H. Create a Gateway Load Balancer (GLB). Configure the Auto Scaling group to register instances with the GLB's target group.

Answer: C

Explanation:

To distribute traffic from customers to EC2 instances in an Auto Scaling group and encrypt all traffic at all stages between the customers and the application servers without decryption at intermediate points, the company should create a Network Load Balancer (NLB) with a TCP listener and configure the Auto Scaling group to register instances with the NLB's target group (Option C). This solution allows for end-to-end encryption of traffic without decryption at intermediate points.

NEW QUESTION 21

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 26

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.
- 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-B.
- G. Advertise the VPC-B CIDR block under the allowed prefixes.
- H. Configure inter-Region transit gateway peering between TGW-A and TGW-B.
- 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 29

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 VPC.
- 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 environment.
- 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 service.
- 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 environment.
- H. Create attachments to the production VPC and the web service VPC.
- 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 32

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 36

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 quota.
- D. Use AWS Site-to-Site VPN to 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 gateway
- 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 transit gateway, 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 39

A company uses a hybrid architecture and has an AWS Direct Connect connection between its on-premises data center and AWS. The company has production applications that run in the on-premises data center. The company also has production applications that run in a VPC. The applications that run in the on-premises data center need to communicate with the applications that run in the VPC. The company is using corp.example.com as the domain name for the on-premises resources and is using an Amazon Route 53 private hosted zone for aws.example.com to host the VPC resources.

The company is using an open-source recursive DNS resolver in a VPC subnet and is using a DNS resolver in the on-premises data center. The company's on-premises DNS resolver has a forwarder that directs requests for the aws.example.com domain name to the DNS resolver in the VPC. The DNS resolver in the VPC has a forwarder that directs requests for the corp.example.com domain name to the DNS resolver in the on-premises data center. The company has decided to replace the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints.

Which combination of steps should a network engineer take to make this replacement? (Choose three.)

- A. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the outbound endpoint.
- B. Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.
- C. Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint.
- D. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.
- E. Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver.
- F. Configure the on-premises DNS resolver to forward aws.example.com queries to the IP addresses of the outbound endpoint.

Answer: BCE

Explanation:

To replace the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints in a hybrid architecture where on-premises applications need to communicate with applications running in a VPC, a network engineer should take the following steps:

- > Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint. (Option C)
- > Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint. (Option B)
- > Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver. (Option E)

These steps will allow for seamless replacement of the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints and enable communication between on-premises and VPC applications.

NEW QUESTION 42

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 47

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 AWS
- 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 deployed
- G. Specify the NAT gateway type as private
- 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 52

A company delivers applications over the internet. An Amazon Route 53 public hosted zone is the authoritative DNS service for the company and its internet applications, all of which are offered from the same domain name.

A network engineer is working on a new version of one of the applications. All the application's components are hosted in the AWS Cloud. The application has a three-tier design. The front end is delivered through Amazon EC2 instances that are deployed in public subnets with Elastic IP addresses assigned. The backend components are deployed in private subnets from RFC1918.

Components of the application need to be able to access other components of the application within the application's VPC by using the same host names as the host names that are used over the public internet. The network engineer also needs to accommodate future DNS changes, such as the introduction of new host names or the retirement of DNS entries.

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

- A. Add a geoproximity routing policy in Route 53.
- B. Create a Route 53 private hosted zone for the same domain name Associate the application's VPC with the new private hosted zone.
- C. Enable DNS hostnames for the application's VPC.
- D. Create entries in the private hosted zone for each name in the public hosted zone by using the corresponding private IP addresses.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs when AWS CloudTrail logs a Route 53 API call to the public hosted zone.
- F. Create an AWS Lambda function as the target of the rule.
- G. Configure the function to use the event information to update the private hosted zone.
- H. Add the private IP addresses in the existing Route 53 public hosted zone.

Answer: BCD

NEW QUESTION 57

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 launching.
- 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 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 MACsec.
- 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 VIF.
- 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 66

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::EC2NetworkInterface"
- 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

NEW QUESTION 71

A company is planning a migration of its critical workloads from an on-premises data center to Amazon EC2 instances. The plan includes a new 10 Gbps AWS Direct Connect dedicated connection from the on-premises data center to a VPC that is attached to a transit gateway. The migration must occur over encrypted paths between the on-premises data center and the AWS Cloud.

Which solution will meet these requirements while providing the HIGHEST throughput?

- A. Configure a public VIF on the Direct Connect connectio
- B. Configure an AWS Site-to-Site VPN connection to the transit gateway as a VPN attachment.
- C. Configure a transit VIF on the Direct Connect connectio
- D. Configure an IPsec VPN connection to an EC2 instance that is running third-party VPN software.
- E. Configure MACsec for the Direct Connect connectio
- F. Configure a transit VIF to a Direct Connect gateway that is associated with the transit gateway.
- G. Configure a public VIF on the Direct Connect connectio
- H. Configure two AWS Site-to-Site VPN connections to the transit gatewa
- I. Enable equal-cost multi-path (ECMP) routing.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/adding-macsec-security-to-aws-direct-connect-c>

NEW QUESTION 76

A company has two AWS accounts one for Production and one for Connectivity. A network engineer needs to connect the Production account VPC to a transit gateway in the Connectivity account. The feature to auto accept shared attachments is not enabled on the transit gateway.

Which set of steps should the network engineer follow in each AWS account to meet these requirements?

- A. * 1. In the Production account: Create a resource share in AWS Resource Access Manager for the transit gatewa
- B. Provide the Connectivity account I
- C. Enable the feature to allow external accounts* 2. In the Connectivity account: Accept the resource.* 3. In the Connectivity account: Create an attachment to the VPC subnets.* 4. In the Production account: Accept the attachmen
- D. Associate a route table with the attachment.
- E. * 1. In the Production account: Create a resource share in AWS Resource Access Manager for the VPC subnet
- F. Provide the Connectivity account I
- G. Enable the feature to allow external accounts.* 2. In the Connectivity account: Accept the resource.* 3. In the Production account: Create an attachment on the transit gateway to the VPC subnets.* 4. In the Connectivity account: Accept the attachmen
- H. Associate a route table with the attachment.
- I. * 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the VPC subnet
- J. Provide the Production account I
- K. Enable the feature to allow external accounts.* 2. In the Production account: Accept the resource.* 3. In the Connectivity account: Create an attachment on the transit gateway to the VPC subnets.* 4. In the Production account: Accept the attachmen
- L. Associate a route table with the attachment.
- M. * 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the transit gatewa
- N. Provide the Production account ID Enable the feature to allow external accounts.* 2. In the Production account: Accept the resource.* 3. In the Production account: Create an attachment to the VPC subnets.* 4. In the Connectivity account: Accept the attachmen
- O. Associate a route table with the attachment.

Answer: A

Explanation:

step 1: In the Production account, create a resource share in AWS Resource Access Manager for the transit gateway and provide the Connectivity account ID. Enabling the feature to allow external accounts is also required to share resources between accounts. Step 2: In the Connectivity account, accept the shared resource. This action will allow the Production account to use the transit gateway in the Connectivity account. Step 3: In the Connectivity account, create an attachment to the VPC subnets. This attachment will enable communication between the VPC in the Production account and the transit gateway in the Connectivity account. Step 4: In the Production account, accept the attachment and associate a route table with the attachment. This will enable the VPC to route traffic through the transit gateway to other resources in the Connectivity account.

NEW QUESTION 78

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 80

A company is hosting an application on Amazon EC2 instances behind a Network Load Balancer (NLB). A solutions architect added EC2 instances in a second Availability Zone to improve the availability of the application. The solutions architect added the instances to the NLB target group.

The company's operations team notices that traffic is being routed only to the instances in the first Availability Zone.

What is the MOST operationally efficient solution to resolve this issue?

- A. Enable the new Availability Zone on the NLB
- B. Create a new NLB for the instances in the second Availability Zone
- C. Enable proxy protocol on the NLB
- D. Create a new target group with the instances in both Availability Zones

Answer: A

Explanation:

When adding instances in a new Availability Zone to an existing Network Load Balancer (NLB), it is important to ensure that the new Availability Zone is enabled on the NLB. This will allow traffic to be routed to instances in both Availability Zones. This can be done by editing the settings of the NLB and selecting the new Availability Zone from the list of available zones.

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 89

A company is building its website on AWS in a single VPC. The VPC has public subnets and private subnets in two Availability Zones. The website has static content such as images. The company is using Amazon S3 to store the content.

The company has deployed a fleet of Amazon EC2 instances as web servers in a private subnet. The EC2 instances are in an Auto Scaling group behind an Application Load Balancer. The EC2 instances will serve traffic, and they must pull content from an S3 bucket to render the webpages. The company is using AWS Direct Connect with a public VIF for on-premises connectivity to the S3 bucket.

A network engineer notices that traffic between the EC2 instances and Amazon S3 is routing through a NAT gateway. As traffic increases, the company's costs are increasing. The network engineer needs to change the connectivity to reduce the NAT gateway costs that result from the traffic between the EC2 instances and Amazon S3.

Which solution will meet these requirements?

- A. Create a Direct Connect private VIF
- B. Migrate the traffic from the public VIF to the private VIF.
- C. Create an AWS Site-to-Site VPN tunnel over the existing public VIF.
- D. Implement interface VPC endpoints for Amazon S3. Update the VPC route table.
- E. Implement gateway VPC endpoints for Amazon S3. Update the VPC route table.

Answer: D

NEW QUESTION 93

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 94

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 VPC.
- 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 disabled.
- H. 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 96

.....

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](#)