CertShared

Amazon

Exam Questions AWS-Certified-Advanced-Networking-Specialty

Amazon AWS Certified Advanced Networking - Specialty

NEW QUESTION 1
A government contractor is designing a multi-account environment with multiple VPCs for a customer. A network security policy requires all traffic between any two VPCs to be transparently inspected by a
third-party appliance.
The customer wants a solution that features AWS Transit Gateway. The setup must be highly available across multiple Availability Zones, and the solution needs to support automated failover. Furthermore, asymmetric routing is not supported by the inspection appliances.
Which combination of steps is part of a solution that meets these requirements? (Choose two.)

A. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VP
B. Connect the inspection VPC to the transit gateway by using a VPCattachmen
C. Create a target group, and register the appliances with the target grou
D. Create a Network Load Balancer (NLB), and set it up to forward to the newly created target grou
E. Configure a default route in the inspection VPCs transit gateway subnet toward the NLB.
F. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VP
G. Connect the inspection VPC to the transit gateway by using a VPC attachmen
H. Create a target group, and register the appliances with the target grou
I. Create a Gateway Load Balancer, and set it up to forward to the newly created target grou
J. Configure a default route in the inspection VPC's transit gateway subnet toward the Gateway Load Balancer endpoint.
K. Configure two route tables on the transit gatewa
L. Associate one route table with all the attachments of the application VPC
M. Associate the other route table with the inspection VPC's attachmen
N. Propagate all VPC attachments into the inspection route tabl
O. Define a static default route in the application route tabl
P. Enable appliance mode on the attachment that connects the inspection VPC.
Q. Configure two route tables on the transit gatewa
R. Associate one route table with all the attachments of the application VPC
S. Associate the other route table with the inspection VPCs attachmen
T. Propagate all VPC attachments into the application route tabl
. Define a static default route in the inspection route tabl
. Enable appliance mode on the attachment that connects the inspection VPC.
. Configure one route table on the transit gatewa
. Associate the route table with all the VPC
. Propagate all VPC attachments into the route tabl
. Define a static default route in the route table.

Answer: BC

NEW QUESTION 2
A company is migrating an existing application to a new AWS account. The company will deploy the application in a single AWS Region by using one VPC and multiple Availability Zones. The application will run on Amazon EC2 instances. Each Availability Zone will have several EC2 instances. The EC2 instances will be deployed in private subnets.
The company's clients will connect to the application by using a web browser with the HTTPS protocol. Inbound connections must be distributed across the Availability Zones and EC2 instances. All connections from the same client session must be connected to the same EC2 instance. The company must provide end-to-end encryption for all connections between the clients and the application by using the application SSL certificate.
Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 3
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 NLUse 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 4
A company has an AWS Site-to-Site VPN connection between its existing VPC and on-premises network. The default DHCP options set is associated with the VPC. The company has an application that is running on an Amazon Linux 2 Amazon EC2 instance in the VPC. The application must retrieve an Amazon RDS database secret that is stored in AWS Secrets Manager through a private VPC endpoint. An on-premises application provides internal RESTful API service that can be reached by URL (https://api.example.internal). Two on-premises Windows DNS servers provide internal DNS resolution.
The application on the EC2 instance needs to call the internal API service that is deployed in the on-premises environment. When the application on the EC2 instance attempts to call the internal API service by referring to the hostname that is assigned to the service, the call fails. When a network engineer tests the API service call from the same EC2 instance by using the API service's IP address, the call is successful.
What should the network engineer do to resolve this issue and prevent the same problem from affecting other resources in the VPC?

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

Answer: B

Explanation:
Creating an Amazon Route 53 Resolver rule and associating it with the VPC would enable forwarding of DNS queries for a specified domain name (example.internal) to a specified IP address (the on-premises Windows DNS servers)3. This would allow EC2 instances in the VPC to resolve the internal API service by using its hostname. Configuring the rule to forward DNS queries only if the domain name matches example.internal would also allow EC2 instances to use the Amazon Route 53 Resolver server for other DNS queries, such as those for AWS services through private VPC endpoints2.

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

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

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

Answer: C

NEW QUESTION 7
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 destinatio
D. Use Amazon Athena to determine which error messages the ALB is receiving.
E. Configure the Amazon S3 bucket destinatio
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 todetermine 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 8
A company has hundreds of VPCs on AWS. All the VPCs access the public endpoints of Amazon S3 and AWS Systems Manager through NAT gateways. All the traffic from the VPCs to Amazon S3 and Systems Manager travels through the NAT gateways. The company's network engineer must centralize access to these services and must eliminate the need to use public endpoints.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a central egress VPC that has private NAT gateway
B. Connect all the VPCs to the central egress VPC by using AWS Transit Gatewa
C. Use the private NAT gateways to connect to Amazon S3 and Systems Manager by using private IP addresses.
D. Create a central shared services VP
E. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
F. Ensure that private DNS is turned of
G. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
H. Create an Amazon Route 53 forwarding rule for each interface VPC endpoin
I. Associate the forwarding rules with all the VPC
J. Forward DNS queries to the interface VPC endpoints in the shared services VPC.
K. Create a central shared services VPIn the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
L. Ensure that private DNS is turned of
M. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
N. Create an Amazon Route 53 private hosted zone with a full service endpoint name for Amazon S3 and Systems Manage
O. Associate the private hosted zones with all the VPC
P. Create an alias record in each private hosted zone with the full AWS service endpoint pointing to the interface VPC endpoint in the shared services VPC.
Q. Create a central shared services VP
R. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
S. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
T. Ensure that private DNS is turned on for the interface VPC endpoints and that the transit gateway is created with DNS support turned on.

Answer: B

Explanation:
Interface VPC endpoints enable private connectivity between VPCs and supported AWS serviceswithout requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection2. Interface VPC endpoints are powered by AWS PrivateLink, a technology that enables private access to AWS

services2. Amazon S3 and AWS Systems Manager support interface VPC endpoin2ts. By turning off private DNS, the interface VPC endpoints can be accessed by using their private IP addresses2. By using Amazon Route 53 forwarding rules, DNS queries can be resolved to the interface VPC endpoints in the shared services VPC3.

NEW QUESTION 9
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 templat
B. Define the primary network interface to be created in one of the private subnet
C. For the second network interface, select one of the public subnet
D. Choose the BYOIP pool ID as the source of public IP addresses.
E. Configure the primary network interface in a private subnet in the launch templat
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 launchin
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 interfac
J. Use the user data option to run a cloud-init script to allocate a second network interface and to associate anElastic 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
A company's network engineer is designing a hybrid DNS solution for an AWS Cloud workload. Individual teams want to manage their own DNS hostnames for their applications in their development environment. The solution must integrate the application-specific hostnames with the centrally managed DNS hostnames from the on-premises network and must provide bidirectional name resolution. The solution also must minimize management overhead.
Which combination of steps should the network engineer take to meet these requirements? (Choose three.)

A. Use an Amazon Route 53 Resolver inbound endpoint.
B. Modify the DHCP options set by setting a custom DNS server value.
C. Use an Amazon Route 53 Resolver outbound endpoint.
D. Create DNS proxy servers.
E. Create Amazon Route 53 private hosted zones.
F. Set up a zone transfer between Amazon Route 53 and the on-premises DNS.

Answer: ABE

NEW QUESTION 10
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 13
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 i
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 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 hasmonitoring 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 23
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 NLUse 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 27
A company has deployed a software-defined WAN (SD-WAN) solution to interconnect all of its offices. The company is migrating workloads to AWS and needs to extend its SD-WAN solution to support connectivity to these workloads.
A network engineer plans to deploy AWS Transit Gateway Connect and two SD-WAN virtual appliances to provide this connectivity. According to company policies, only a single SD-WAN virtual appliance can handle traffic from AWS workloads at a given time.
How should the network engineer configure routing to meet these requirements?

A. Add a static default route in the transit gateway route table to point to the secondary SD-WAN virtual applianc
B. Add routes that are more specific to point to the primary SD-WAN virtual appliance.
C. Configure the BGP community tag 7224:7300 on the primary SD-WAN virtual appliance for BGP routes toward the transit gateway.
D. Configure the AS_PATH prepend attribute on the secondary SD-WAN virtual appliance for BGP routes toward the transit gateway.
E. Disable equal-cost multi-path (ECMP) routing on the transit gateway for Transit Gateway Connect.

Answer: A

NEW QUESTION 29
A company is using Amazon Route 53 Resolver DNS Firewall in a VPC to block all domains except domains that are on an approved list. The company is concerned that if DNS Firewall is unresponsive, resources in the VPC might be affected if the network cannot resolve any DNS queries. To maintain application service level agreements, the company needs DNS queries to continue to resolve even if Route 53 Resolver does not receive a response from DNS Firewall.
Which change should a network engineer implement to meet these requirements?

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

Answer: B

NEW QUESTION 33
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 configuratio
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 configuratio
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 configuratio
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 configuratio
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 rules3. Creating an AWS Systems Manager Automation runbook to remediate noncompliant security groups would enable automation of the remediation process2. 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 36
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 acceleratio
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 gatewa
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 40
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 VPRemove the existing VPN connections that are attached directly to the virtual private gateways.
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 43
A company has a hybrid cloud environment. The company's data center is connected to the AWS Cloud by an AWS Direct Connect connection. The AWS environment includes VPCs that are connected together in a hub-and-spoke model by a transit gateway. The AWS environment has a transit VIF with a Direct Connect gateway for on-premises connectivity.
The company has a hybrid DNS model. The company has configured Amazon Route 53 Resolver endpoints in the hub VPC to allow bidirectional DNS traffic flow.
The company is running a backend application in one of the VPCs.
The company uses a message-oriented architecture and employs Amazon Simple Queue Service (Amazon SQS) to receive messages from other applications over a private network. A network engineer wants to use an interface VPC endpoint for Amazon SQS for this architecture. Client services must be able to access the endpoint service from on premises and from multiple VPCs within the company's AWS infrastructure.
Which combination of steps should the network engineer take to ensure that the client applications can resolve DNS for the interface endpoint? (Choose three.)

A. Create the interface endpoint for Amazon SQS with the option for private DNS names turned on.
B. Create the interface endpoint for Amazon SQS with the option for private DNS names turned off.
C. Manually create a private hosted zone for sqs.us-east-1.amazonaws.co
D. Add necessary records that point to the interface endpoin
E. Associate the private hosted zones with other VPCs.
F. Use the automatically created private hosted zone for sqs.us-east-1.amazonaws.com with previously created necessary records that point to the interface endpoin
G. Associate the private hosted zones with other VPCs.
H. Access the SQS endpoint by using the public DNS name sqs.us-east-1 amazonaws.com in VPCs and on premises.
I. Access the SQS endpoint by using the private DNS name of the interface endpoint.sqs.us-east-1.vpce.amazonaws.com in VPCs and on premises.

Answer: ADF

NEW QUESTION 47
A company has a global network and is using transit gateways to connect AWS Regions together. The company finds that two Amazon EC2 instances in different Regions are unable to communicate with each other. A network engineer needs to troubleshoot this connectivity issue.
What should the network engineer do to meet this requirement?

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

Answer: C

Explanation:
Using AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between VPCs and transit gateways1. Verifying that the VPC route tables are correct would enable identification of routing issues within a VPC. Using VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC would enable identification of traffic filtering issues within a VPC2. Additionally, using VPC Reachability Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between transit gateways in different Regions. VPC Reachability Analyzer is a configuration analysis tool that enables connectivity testing between a source resource and a destination resource in your VPCs.

NEW QUESTION 49
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 51
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 failove
B. Create an origin group for each Region where the solution is deployed.
C. Set up Route 53 latency-based routin
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 Accelerato
G. Configure Regional endpoint groups andhealth checks.
H. Set up Bring Your Own IP (BYOIP) addresse
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 53
A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin in an Amazon CloudFront distribution. The company wants to implement a custom authentication system that will provide a token for its authenticated customers.
The web application must ensure that the GET/POST requests come from authenticated customers before it delivers the content. A network engineer must design a solution that gives the web application the ability to identify authorized customers.
What is the MOST operationally efficient solution that meets these requirements?

A. Use the ALB to inspect the authorized token inside the GET/POST request payloa
B. Use an AWS Lambda function to insert a customized header to inform the web application of an authenticated customer request.
C. Integrate AWS WAF with the ALB to inspect the authorized token inside the GET/POST request payloa
D. Configure the ALB listener to insert a customized header to inform the web application of an authenticated customer request.
E. Use an AWS Lambda@Edge function to inspect the authorized token inside the GET/POST request payloa
F. Use the Lambda@Edge function also to insert a customized header to inform the web application of an authenticated customer request.
G. Set up an EC2 instance that has a third-party packet inspection tool to inspect the authorized token inside the GET/POST request payloa
H. Configure the tool to insert a customized header to inform the web application of an authenticated customer request.

Answer: C

NEW QUESTION 56
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 targe
B. Behind the NL
C. deploy a fleet of EC2 instances in an Auto Scaling grou
D. Use Traffic Mirroring as necessary.
E. Deploy an Application Load Balancer (ALB) as the traffic mirror targe
F. Behind the ALB, deploy a fleet of EC2 instances in an Auto Scaling grou
G. Use Traffic Mirroring only during non-business hours.
H. Deploy a Gateway Load Balancer (GLB) as the traffic mirror targe
I. Behind the GL
J. deploy a fleet of EC2 instances in an Auto Scaling grou
K. Use Traffic Mirroring as necessary.
L. Deploy an Application Load Balancer (ALB) with an HTTPS listener as the traffic mirror targe
M. Behind the AL
N. deploy a fleet of EC2 instances in an Auto Scaling grou
O. Use Traffic Mirroring only during active events or business hours.

Answer: A

NEW QUESTION 58
A Network Engineer is provisioning a subnet for a load balancer that will sit in front of a fleet of application servers in a private subnet. There is limited IP space left in the VPC CIDR. The application has few users now but is expected to grow quickly to millions of users.
What design will use the LEAST amount of IP space, while allowing for this growth?

A. Use two /29 subnets for an Application Load Balancer in different Availability Zones.
B. Use one /29 subnet for the Network Load Balance
C. Add another VPC CIDR to the VPC to allow for future growth.
D. Use two /28 subnets for a Network Load Balancer in different Availability Zones.
E. Use one /28 subnet for an Application Load Balance
F. Add another VPC CIDR to the VPC to allow for future growth.

Answer: C

NEW QUESTION 62
A company has deployed a new web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Amazon EC2 Auto Scaling group. Enterprise customers from around the world will use the application. Employees of these enterprise customers will connect to the application over HTTPS from office locations.
The company must configure firewalls to allow outbound traffic to only approved IP addresses. The employees of the enterprise customers must be able to access the application with the least amount of latency.
Which change should a network engineer make in the infrastructure to meet these requirements?

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

Answer: B

Explanation:

Amazon CloudFront is a content delivery network (CDN) that can speed up the delivery of static and dynamic web content, such as images, videos, and APIs2. CloudFront can also provide end-to-end encryption for HTTPS traffic by using SSL certificates from AWS Certificate Manager (ACM) or other sources2. CloudFron can also support session affinity (sticky sessions) with a load balancer-generated cookie or an application-based cookie policy2.

NEW QUESTION 64
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 66
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 Zonesafter 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 71
A company is deploying a new application on AWS. The application uses dynamic multicasting. The company has five VPCs that are all attached to a transit gateway Amazon EC2 instances in each VPC need to be able to register dynamically to receive a multicast transmission.
How should a network engineer configure the AWS resources to meet these requirements?

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

Answer: C

NEW QUESTION 76
A media company is implementing a news website for a global audience. The website uses Amazon CloudFront as its content delivery network. The backend runs on Amazon EC2 Windows instances behind an Application Load Balancer (ALB). The instances are part of an Auto Scaling group. The company's customers access the website by using service example com as the CloudFront custom domain name. The CloudFront origin points to an ALB that uses service-alb.example.com as the domain name.
The company's security policy requires the traffic to be encrypted in transit at all times between the users and the backend.
Which combination of changes must the company make to meet this security requirement? (Choose three.)

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

Answer: BDE

NEW QUESTION 77
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 deckled 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 78
A security team is performing an audit of a company's AWS deployment. The security team is concerned that two applications might be accessing resources that should be blocked by network ACLs and security groups. The applications are deployed across two Amazon Elastic Kubernetes Service (Amazon EKS) clusters that use the Amazon VPC Container Network Interface (CNI) plugin for Kubernetes. The clusters are in separate subnets within the same VPC and have a Cluster Autoscaler configured.
The security team needs to determine which POD IP addresses are communicating with which services throughout the VPC. The security team wants to limit the number of flow logs and wants to examine the traffic from only the two applications.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create VPC flow logs in the default forma
B. Create a filter to gather flow logs only from the EKS nodes.Include the srcaddr field and the dstaddr field in the flow logs.
C. Create VPC flow logs in a custom forma
D. Set the EKS nodes as the resource Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.
E. Create VPC flow logs in a custom forma
F. Set the application subnets as resource
G. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.
H. Create VPC flow logs in a custom forma
I. Create a filter to gather flow logs only from the EKS nodes.Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.

Answer: D

NEW QUESTION 81
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 85
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Certified-Advanced-Networking-Specialty Practice Exam Features:

* AWS-Certified-Advanced-Networking-Specialty Questions and Answers Updated Frequently

* AWS-Certified-Advanced-Networking-Specialty Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Advanced-Networking-Specialty Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Advanced-Networking-Specialty Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Advanced-Networking-Specialty Practice Test Here](#)