Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam

https://www.2passeasy.com/dumps/ANS-C01/

NEW QUESTION 1
A company is using custom DNS servers that run BIND for name resolution in its VPCs. The VPCs are deployed across multiple AWS accounts that are part of the same organization in AWS Organizations. All the VPCs are connected to a transit gateway. The BIND servers are running in a central VPC and are configured to forward all queries for an on-premises DNS domain to DNS servers that are hosted in an on-premises data center. To ensure that all the VPCs use the custom DNS servers, a network engineer has configured a VPC DHCP options set in all the VPCs that specifies the custom DNS servers to be used as domain name servers.
Multiple development teams in the company want to use Amazon Elastic File System (Amazon EFS). A development team has created a new EFS file system but cannot mount the file system to one of its Amazon EC2 instances. The network engineer discovers that the EC2 instance cannot resolve the IP address for the EFS mount point fs-33444567d.efs.us-east-1.amazonaws.com. The network engineer needs to implement a solution so that development teams throughout the organization can mount EFS file systems.
Which combination of steps will meet these requirements? (Choose two.)

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

Answer: BD

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

NEW QUESTION 2
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 3
A company has developed an application on AWS that will track inventory levels of vending machines and initiate the restocking process automatically. The company plans to integrate this application with vending machines and deploy the vending machines in several markets around the world. The application resides in a VPC in the us-east-1 Region. The application consists of an Amazon Elastic Container Service (Amazon ECS) cluster behind an Application Load Balancer (ALB). The communication from the vending machines to the application happens over HTTPS.
The company is planning to use an AWS Global Accelerator accelerator and configure static IP addresses of the accelerator in the vending machines for application endpoint access. The application must be accessible only through the accelerator and not through a direct connection over the internet to the ALB endpoint.
Which solution will meet these requirements?

A. Configure the ALB in a private subnet of the VP
B. Attach an internet gateway without adding routes in the subnet route tables to point to the internet gatewa
C. Configure the accelerator with endpoint groups that include the ALB endpoin
D. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
E. Configure the ALB in a private subnet of the VP
F. Configure the accelerator with endpoint groups that include the ALB endpoin

G. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
H. Configure the ALB in a public subnet of the VPAttach an internet gatewa
I. Add routes in the subnet route tables to point to the internet gatewa
J. Configure the accelerator with endpoint groups that include the ALB endpoin
K. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.
L. Configure the ALB in a private subnet of the VP
M. Attach an internet gatewa
N. Add routes in the subnet route tables to point to the internet gatewa
O. Configure the accelerator with endpoint groups that include the ALB endpoin
P. Configure the ALB's security group to only allow inbound trafficfrom the accelerator's IP addresses on the ALB listener port.

Answer: A

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

NEW QUESTION 4
A network engineer needs to standardize a company's approach to centralizing and managing interface VPC endpoints for private communication with AWS services. The company uses AWS Transit Gateway for inter-VPC connectivity between AWS accounts through a hub-and-spoke model. The company's network services team must manage all Amazon Route 53 zones and interface endpoints within a shared services AWS account. The company wants to use thiscentralized model to provide AWS resources with access to AWS Key Management Service (AWS KMS) without sending traffic over the public internet.
What should the network engineer do to meet these requirements?

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

Answer: A

NEW QUESTION 5
A company has its production VPC (VPC-A) in the eu-west-1 Region in Account 1. VPC-A is attached to a transit gateway (TGW-A) that is connected to an on-premises data center in Dublin, Ireland, by an AWS
Direct Connect transit VIF that is configured for an AWS Direct Connect gateway. The company also has a staging VPC (VPC-B) that is attached to another transit gateway (TGW-B) in the eu-west-2 Region in Account 2.
A network engineer must implement connectivity between VPC-B and the on-premises data center in Dublin. Which solutions will meet these requirements? (Choose two.)

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

Answer: BC

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

NEW QUESTION 6
A company has deployed a web application on AWS. The web application uses an Application Load Balancer (ALB) across multiple Availability Zones. The targets of the ALB are AWS Lambda functions. The web application also uses Amazon CloudWatch metrics for monitoring.
Users report that parts of the web application are not loading properly. A network engineer needs to troubleshoot the problem. The network engineer enables access logging for the ALB.
What should the network engineer do next to determine which errors the ALB is receiving?

A. Send the logs to Amazon CloudWatch Log
B. Review the ALB logs in CloudWatch Insights to determine which error messages the ALB is receiving.
C. Configure the Amazon S3 bucket 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 7
All IP addresses within a 10.0.0.0/16 VPC are fully utilized with application servers across two Availability Zones. The application servers need to send frequent UDP probes to a single central authentication server on the Internet to confirm that is running up-to-date packages. The network is designed for application servers to use a single NAT gateway for internal access. Testing reveals that a few of the servers are unable to communicate with the authentication server.

A. The NAT gateway does not support UDP traffic.
B. The authentication server is not accepting traffic.
C. The NAT gateway cannot allocate more ports.
D. The NAT gateway is launched in a private subnet.

Answer: C

Explanation:
Ref:https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html
"A NAT gateway can support up to 55,000 simultaneous connections to each unique destination. This limit also applies if you create approximately 900 connections per second to a single destination (about 55,000 connections per minute). If the destination IP address, the destination port, or the protocol (TCP/UDP/ICMP) changes, you can create an additional 55,000 connections. For more than 55,000 connections, there is an increased chance of connection errors due to port allocation errors. These errors can be monitored by viewing the ErrorPortAllocation CloudWatch metric for your NAT gateway. For more information, see Monitoring NAT Gateways Using Amazon CloudWatch."

NEW QUESTION 8
A company is deploying third-party firewall appliances for traffic inspection and NAT capabilities in its VPC. The VPC is configured with private subnets and public subnets. The company needs to deploy the firewall appliances behind a load balancer.
Which architecture will meet these requirements MOST cost-effectively?

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

Answer: B

NEW QUESTION 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 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 10
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 VI
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 13
A company has expanded its network to the AWS Cloud by using a hybrid architecture with multiple AWS accounts. The company has set up a shared AWS account for the connection to its on-premises data centers and the company offices. The workloads consist of private web-based services for internal use. These services run in different AWS accounts. Office-based employees consume these services by using a DNS name in an on-premises DNS zone that is named example.internal.
The process to register a new service that runs on AWS requires a manual and complicated change request to the internal DNS. The process involves many teams.
The company wants to update the DNS registration process by giving the service creators access that will allow them to register their DNS records. A network engineer must design a solution that will achieve this goal. The solution must maximize cost-effectiveness and must require the least possible number of configuration changes.
Which combination of steps should the network engineer take to meet these requirements? (Choose three.)

A. Create a record for each service in its local private hosted zone (serviceA.account1.aws.example.internal). Provide this DNS record to the employees who need access.
B. Create an Amazon Route 53 Resolver inbound endpoint in the shared account VP
C. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS server
D. Set the forwarding IP addresses to the inbound endpoint's IP addresses that were created.
E. Create an Amazon Route 53 Resolver rule to forward any queries made to onprem.example.internal to the on-premises DNS servers.
F. Create an Amazon Route 53 private hosted zone named aws.example.internal in the shared AWSaccount to resolve queries for this domain.
G. Launch two Amazon EC2 instances in the shared AWS accoun
H. Install BIND on each instanc
I. Create a DNS conditional forwarder on each BIND server to forward queries for each subdomain under aws.example.internal to the appropriate private hosted zone in each AWS accoun
J. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS server
K. Set the forwarding IP addresses to the IP addresses of the BIND servers.
L. Create a private hosted zone in the shared AWS account for each account that runs the service.Configure the private hosted zone to contain aws.example.internal in the domain (account1.aws.example.internal). Associate the private hosted zone with the VPC that runs the service and the shared account VPC.

Answer: ABD

Explanation:
To meet the requirements of updating the DNS registration process while maximizing cost-effectiveness and minimizing configuration changes, the network engineer should take the following steps:

≫ Create an Amazon Route 53 Resolver inbound endpoint in the shared account VPC. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS servers. Set the forwarding IP addresses to the inbound endpoint's IP addresses that were created (Option B).

≫ Create an Amazon Route 53 private hosted zone named aws.example.internal in the shared AWS account to resolve queries for this domain (Option D).

≫ Create a record for each service in its local private hosted zone (serviceA.account1.aws.example.internal). Provide this DNS record to the employees who need access (Option A).

These steps will allow service creators to register their DNS records while keeping costs low and minimizing configuration changes.

NEW QUESTION 15
A company manages resources across VPCs in multiple AWS Regions. The company needs to connect to the resources by using its internal domain name. A network engineer needs to apply the aws.example.com DNS suffix to all resources.
What must the network engineer do to meet this requirement?

A. Create an Amazon Route 53 private hosted zone for aws.example.com in each Region that has resource
B. Associate the private hosted zone with that Region's VP
C. In the appropriate private hosted zone, create DNS records for the resources in each Region.
D. Create one Amazon Route 53 private hosted zone for aws.example.co
E. Configure the private hosted zone to allow zone transfers with every VPC.
F. Create one Amazon Route 53 private hosted zone for example.co
G. Create a single resource record for aws.example.com in the private hosted zon
H. Apply a multivalue answer routing policy to the recor
I. Add all VPC resources as separate values in the routing policy.
J. Create one Amazon Route 53 private hosted zone for aws.example.co
K. Associate the private hosted zone with every VPC that has resource
L. In the private hosted zone, create DNS records for all resources.

Answer: D

Explanation:
Creating one private hosted zone for aws.example.com and associating it with every VPC that has resources would enable DNS resolution for all resources by using their internal domain name. Creating 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 would enable private connectivity to Amazon S3 and AWS Systems Manager without using public endpoints.

NEW QUESTION 17
A company is using a NAT gateway to allow internet connectivity for private subnets in a VPC in the us-west-2 Region. After a security audit, the company needs to remove the NAT gateway.
In the private subnets, the company has resources that use the unified Amazon CloudWatch agent. A network engineer must create a solution to ensure that the unified CloudWatch agent continues to work after the removal of the NAT gateway.
Which combination of steps should the network engineer take to meet these requirements? (Choose three.)

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

Answer: BDF

NEW QUESTION 20
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 22
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 26
A company has deployed an AWS Network Firewall firewall into a VPC. A network engineer needs to implement a solution to deliver Network Firewall flow logs to the company's Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster in the shortest possible time.
Which solution will meet these requirements?

A. Create an Amazon S3 bucke
B. Create an AWS Lambda function to load logs into the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluste
C. Enable Amazon Simple Notification Service (Amazon SNS) notifications on the S3 bucket to invoke the Lambda functio
D. Configure flow logs for the firewall
E. Set the S3 bucket as the destination.
F. Create an Amazon Kinesis Data Firehose delivery stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destinatio
G. Configure flow logs for the firewall Set the Kinesis Data Firehose delivery stream as the destination for the Network Firewall flow logs.
H. Configure flow logs for the firewall
I. Set the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination for the Network Firewall flow logs.
J. Create an Amazon Kinesis data stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destinatio
K. Configure flow logs for the firewall
L. Set the Kinesis data stream as the destination for the Network Firewall flow logs.

Answer: B

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-analyze-aws-network-firewall-logs-usin

NEW QUESTION 30
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 31
A company is planning to create a service that requires encryption in transit. The traffic must not be decrypted between the client and the backend of the service. The company will implement the service by using the gRPC protocol over TCP port 443. The service will scale up to thousands of simultaneous connections. The backend of the service will be hosted on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with the Kubernetes Cluster Autoscaler and the Horizontal Pod Autoscaler configured. The company needs to use mutual TLS for two-way authentication between the client and the backend.
Which solution will meet these requirements?

A. Install the AWS Load Balancer Controller for Kubernete
B. Using that controller, configure a NetworkLoad Balancer with a TCP listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
C. Install the AWS Load Balancer Controller for Kubernete
D. Using that controller, configure an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the IP addresses of the backend service

Pods.
E. Create a target grou
F. Add the EKS managed node group's Auto Scaling group as a target Create an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the target group.
G. Create a target grou
H. Add the EKS managed node group's Auto Scaling group as a targe
I. Create a Network Load Balancer with a TLS listener on port 443 to forward traffic to the target group.

Answer: B

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

NEW QUESTION 34
A company has deployed a critical application on a fleet of Amazon EC2 instances behind an Application Load Balancer. The application must always be reachable on port 443 from the public internet. The application recently had an outage that resulted from an incorrect change to the EC2 security group.
A network engineer needs to automate a way to verify the network connectivity between the public internet and the EC2 instances whenever a change is made to the security group. The solution also must notify the network engineer when the change affects the connection.
Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 36
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 40
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 MACse
E. Configure MACsec on the edge router

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

Answer: C

NEW QUESTION 44
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 45
A company is planning to deploy many software-defined WAN (SD-WAN) sites. The company is using AWS Transit Gateway and has deployed a transit gateway in the required AWS Region. A network engineer needs to deploy the SD-WAN hub virtual appliance into a VPC that is connected to the transit gateway. The solution must support at least 5 Gbps of throughput from the SD-WAN hub virtual appliance to other VPCs that are attached to the transit gateway.
Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 48
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 52
A company is deploying a new application in the AWS Cloud. The company wants a highly available web server that will sit behind an Elastic Load Balancer. The load balancer will route requests to multiple target groups based on the URL in the request. All traffic must use HTTPS. TLS processing must be offloaded to the load balancer. The web server must know the user's IP address so that the company can keep accurate logs for security purposes.

Which solution will meet these requirements?

A. Deploy an Application Load Balancer with an HTTPS listene
B. Use path-based routing rules to forward the traffic to the correct target grou
C. Include the X-Forwarded-For request header with traffic to the targets.
D. Deploy an Application Load Balancer with an HTTPS listener for each domai
E. Use host-based routing rules to forward the traffic to the correct target group for each domai
F. Include the X-Forwarded-For request header with traffic to the targets.
G. Deploy a Network Load Balancer with a TLS listene
H. Use path-based routing rules to forward the traffic to the correct target grou
I. Configure client IP address preservation for traffic to the targets.
J. Deploy a Network Load Balancer with a TLS listener for each domai
K. Use host-based routing rules to forward the traffic to the correct target group for each domai
L. Configure client IP address preservation for traffic to the targets.

Answer: A

Explanation:
An Application Load Balancer (ALB) can be used to route traffic to multiple target groups based on the URL in the request. The ALB can be configured with an HTTPS listener to ensure all traffic uses HTTPS. TLS processing can be offloaded to the ALB, which reduces the load on the web server. Path-based routing rules can be used to route traffic to the correct target group based on the URL in the request. The X-Forwarded-For request header can be included with traffic to the targets, which will allow the web server to know the user's IP address and keep accurate logs for security purposes.

NEW QUESTION 53
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 54
A company has been using an outdated application layer protocol for communication among applications. The company decides not to use this protocol anymore and must migrate all applications to support a new protocol. The old protocol and the new protocol are TCP-based, but the protocols use different port numbers.
After several months of work, the company has migrated dozens of applications that run on Amazon EC2 instances and in containers. The company believes that all the applications have been migrated, but the company wants to verify this belief. A network engineer needs to verify that no application is still using the old protocol.
Which solution will meet these requirements without causing any downtime?

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

Answer: C

Explanation:
Configuring VPC flow logs to be delivered into an Amazon S3 bucket would enable capture of information about the IP traffic going to and from network interfaces within the VPC3. Using Amazon Athena to query the data and to filter for the port number that is used by the old protocol would enable identification of applications that are still using the old protocol.

NEW QUESTION 55
A network engineer is designing a hybrid architecture that uses a 1 Gbps AWS Direct Connect connection between the company's data center and two AWS Regions: us-east-1 and eu-west-1. The VPCs in us-east-1 are connected by a transit gateway and need to access several on-premises databases. According to company policy, only one VPC in eu-west-1 can be connected to one on-premises server. The on-premises network segments the traffic between the databases and the server.
How should the network engineer set up the Direct Connect connection to meet these requirements?

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

Answer: B

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

NEW QUESTION 60
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 observed.Create a new 10 Gbps dedicated connectio
I. 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 62
A bank built a new version of its banking application in AWS using containers that content to an on-premises database over VPN connection. This application version requires users to also update their client application. The bank plans to deprecate the earlier client version. However, the company wants to keep supporting earlier clients through their on-premises version of the application to serve a small portion of the customers who haven't yet upgraded.
What design will allow the company to serve both newer and earlier clients in the MOST efficient way?

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

Answer: D

NEW QUESTION 65
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 70
A company has several production applications across different accounts in the AWS Cloud. The company operates from the us-east-1 Region only. Only certain partner companies can access the applications. The applications are running on Amazon EC2 instances that are in an Auto Scaling group behind an Application Load Balancer (ALB). The EC2 instances are in private subnets and allow traffic only from the ALB. The ALB is in a public subnet and allows inbound traffic only from partner network IP address ranges over port 80.
When the company adds a new partner, the company must allow the IP address range of the partner network in the security group that is associated with the ALB in each account. A network engineer must implement a solution to centrally manage the partner network IP address ranges.
Which solution will meet these requirements in the MOST operationally efficient manner?

A. Create an Amazon DynamoDB table to maintain all IP address ranges and security groups that need to be update
B. Update the DynamoDB table with the new IP address range when the company adds a new partne
C. Invoke an AWS Lambda function to read new IP address ranges and security groups from the DynamoDB table to update the security group
D. Deploy this solution in all accounts.
E. Create a new prefix lis
F. Add all allowed IP address ranges to the prefix lis
G. Use Amazon EventBridge (Amazon CloudWatch Events) rules to invoke an AWS Lambda function to update security groups whenever a new IP address range is added to the prefix lis
H. Deploy this solution in all accounts.
I. Create a new prefix lis
J. Add all allowed IP address ranges to the prefix lis
K. Share the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM). Update security groups to use the prefix list instead of the partner IP address rang
L. Update the prefix list with the new IP address range when the company adds a new partner.
M. Create an Amazon S3 bucket to maintain all IP address ranges and security groups that need to be update
N. Update the S3 bucket with the new IP address range when the company adds a new partne
O. Invoke an AWS Lambda function to read new IP address ranges and security groups from the S3 bucket to update the security group
P. Deploy this solution in all accounts.

Answer: C

Explanation:
Creating a new prefix list and adding all allowed IP address ranges to the prefix list would enable grouping of CIDR blocks that can be referenced in security group rules3. Sharing the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM)would enable central management of the partner network IP address ranges5. Updating security groups to use the prefix list instead of the partner IP address range would enable simplification of security group rules3. Updating the prefix list with the new IP address range when the company adds a new partner would enable automatic propagation of the changes to all security groups that use the prefix list3.

NEW QUESTION 73
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 78
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual ANS-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the ANS-C01 Product From:

https://www.2passeasy.com/dumps/ANS-C01/

Money Back Guarantee

ANS-C01 Practice Exam Features:

* ANS-C01 Questions and Answers Updated Frequently

* ANS-C01 Practice Questions Verified by Expert Senior Certified Staff

* ANS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* ANS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year