



Amazon-Web-Services

Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam

NEW QUESTION 1

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

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

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

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

Answer: BD

Explanation:

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

NEW QUESTION 2

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

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

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

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

Answer: B

NEW QUESTION 3

A company has 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 4

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 5

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/vpn/latest/s2svpn/private-ip-dx.html>

An IPsec VPN connection over the transit VIF can encrypt traffic between the on-premises network and AWS without using public IP addresses or the internet². A VPC endpoint for Amazon S3 can enable private access to S3 buckets within the same region.HTTPS can provide additional encryption for communication.

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 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 Gateway
- 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 VPC
- E. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access
- F. Ensure that private DNS is turned off
- G. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway
- H. Create an Amazon Route 53 forwarding rule for each interface VPC endpoint
- 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 VPC. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access
- L. Ensure that private DNS is turned off
- M. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway
- N. Create an Amazon Route 53 private hosted zone with a full service endpoint name for Amazon S3 and Systems Manager
- 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 VPC
- R. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access
- S. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway
- 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 services without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Interface VPC endpoints are powered by AWS PrivateLink, a technology that enables private access to AWS services. Amazon S3 and AWS Systems Manager support interface VPC endpoints. By turning off private DNS, the interface VPC endpoints can be accessed by using their private IP addresses. By using Amazon Route 53 forwarding rules, DNS queries can be resolved to the interface VPC endpoints in the shared services VPC.

NEW QUESTION 8

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 it 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 9

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

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

A recent design meeting revealed that the customers have IP address overlap with the provider's AWS deployment. The customers have stated that they will not

share their internal IP addresses and that they do not want to connect to the provider's SaaS service over the internet. Which combination of steps is part of a solution that meets these requirements? (Choose two.)

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

Answer: AB

Explanation:

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

NEW QUESTION 10

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

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

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

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 11

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 connection
- 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 connection
- D. Configure an IPsec VPN connection to an EC2 instance that is running third-party VPN software.
- E. Configure MACsec for the Direct Connect connection
- 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 connection
- H. Configure two AWS Site-to-Site VPN connections to the transit gateway
- 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 14

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

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

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

- A. Set up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional traffic load from remote employees and the additional application
- B. Create a link aggregation group (LAG).
- C. Deploy an AWS Site-to-Site VPN connection to the application VPC
- D. Configure the on-premises routing for the remote employees to connect to the Site-to-Site VPN connection.
- E. Deploy Amazon Workspaces into the application VPC. Instruct the remote employees to connect to Workspaces.

- F. Replace the existing 1 Gbps Direct Connect connection with two new 2 Gbps Direct Connect hosted connection
- G. Create an AWS Client VPN endpoint in the application VP
- H. Instruct the remote employees to connect to the Client VPN endpoint.

Answer: A

Explanation:

Setting up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional trafficload from remote employees and the additional application would provide more bandwidth and lower latency than a VPN connection over the public internet¹. Creating a link aggregation group (LAG) with the existing and new Direct Connect connections would provide resiliency and redundancy for the AWS connectivity².

NEW QUESTION 17

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

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

What should a network engineer do to reduce this latency?

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

Answer: A

Explanation:

Enabling acceleration for a Site-to-Site VPN connection uses AWS Global Accelerator to route traffic from the on-premises network to an AWS edge location that is closest to the customer gateway device¹. AWS Global Accelerator optimizes the network path, using the congestion-free AWS global network to route traffic to the endpoint that provides the best application performance². Setting the transit gateway as the target gateway enables connectivity between the on-premises network and multiple VPCs that are attached to the transit gateway³.

NEW QUESTION 20

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 22

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

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

AWSTemplateFormation Version: 2010-09-09 Parameters:

Originating VPCId: Type: String RemoteVPCId: Type: String

RemoteVPCAccountId: Type: String Resources:

newVPCPeeringConnection:

Type: 'AWS::EC2::VPCPeeringConnection'

Properties:

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

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

- A. Resources:NewEC2SecurityGroup:Type: AWS::EC2::SecurityGroup
- B. Resources:NetworkInterfaceToRemoteVPC:Type: "AWS::EC2NetworkInterface"
- C. Resources:newEC2Route:Type: AWS::EC2::Route

- D. Resources:VPCGatewayToRemoteVPC:Type: "AWS::EC2::VPCGatewayAttachment"
E. Resources:newVPCPeeringConnection:Type: 'AWS::EC2VPCPeeringConnection'PeerRoleArn: !Ref PeerRoleArn

Answer: CE

Explanation:

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/AWS_EC2.html

NEW QUESTION 25

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 27

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 29

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 31

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

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

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

Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 33

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 bucket
- B. Create an AWS Lambda function to load logs into the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster
- C. Enable Amazon Simple Notification Service (Amazon SNS) notifications on the S3 bucket to invoke the Lambda function
- 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 destination
- 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 destination
- 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-using-aws-lambda/>

NEW QUESTION 38

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

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

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

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

Answer: C

NEW QUESTION 41

A banking company is successfully operating its public mobile banking stack on AWS. The mobile banking stack is deployed in a VPC that includes private subnets

and public subnets. The company is using IPv4 networking and has not deployed or supported IPv6 in the environment. The company has decided to adopt a third-party service provider's API and must integrate the API with the existing environment. The service provider's API requires the use of IPv6.

A network engineer must turn on IPv6 connectivity for the existing workload that is deployed in a private subnet. The company does not want to permit IPv6 traffic from the public internet and mandates that the company's servers must initiate all IPv6 connectivity. The network engineer turns on IPv6 in the VPC and in the private subnets.

Which solution will meet these requirements?

- A. Create an internet gateway and a NAT gateway in the VP
- B. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT gateway.
- C. Create an internet gateway and a NAT instance in the VP
- D. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT instance.
- E. Create an egress-only Internet gateway in the VPAdd a route to the existing subnet route tables to point IPv6 traffic to the egress-only internet gateway.
- F. Create an egress-only internet gateway in the VP
- G. Configure a security group that denies all inbound traffic
- H. Associate the security group with the egress-only internet gateway.

Answer: C

NEW QUESTION 43

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 48

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 49

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

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

What should a network engineer do to resolve this issue?

- A. Modify the ALB listener configuration
- B. Edit the rule that forwards traffic to the target group
- C. Change the rule to enable group-level stickiness
- D. Set the duration to the maximum application session length.
- E. Replace the ALB with a Network Load Balance
- F. Create a TLS listener
- G. Create a new target group with the protocol type set to TLS Register the EC2 instance
- H. Modify the target group configuration by enabling the stickiness attribute.
- I. Modify the ALB target group configuration by enabling the stickiness attribute
- J. Use an application-based cookie

- K. Set the duration to the maximum application session length.
- L. Remove the 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 51

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 54

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 amountof 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 55

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 57

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

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

Answer: C

NEW QUESTION 62

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

NEW QUESTION 68

A company is migrating an application from on premises to AWS. The company will host the application on Amazon EC2 instances that are deployed in a single VPC. During the migration period, DNS queries from the EC2 instances must be able to resolve names of on-premises servers. The migration is expected to take 3 months. After the 3-month migration period, the resolution of on-premises servers will no longer be needed. What should a network engineer do to meet these requirements with the LEAST amount of configuration?

- A. Set up an AWS Site-to-Site VPN connection between on premises and AWS
- B. Deploy an Amazon Route 53 Resolver outbound endpoint in the Region that is hosting the VPC.

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

Answer: A

Explanation:

Setting up an AWS Site-to-Site VPN connection between on premises and AWS would enable a secure and encrypted connection over the public internet¹. Deploying an Amazon Route 53 Resolver outbound endpoint in the Region that is hosting the VPC would enable forwarding of DNS queries for on-premises servers to the on-premises DNS servers². This would allow EC2 instances in the VPC to resolve names of on-premises servers during the migration period. After the migration period, the Route 53 Resolver outbound endpoint can be deleted with minimal configuration changes.

NEW QUESTION 69

A company delivers applications over the internet. An Amazon Route 53 public hosted zone is the authoritative DNS service for the company and its internet applications, all of which are offered from the same domain name. A network engineer is working on a new version of one of the applications. All the application's components are hosted in the AWS Cloud. The application has a three-tier design. The front end is delivered through Amazon EC2 instances that are deployed in public subnets with Elastic IP addresses assigned. The backend components are deployed in private subnets from RFC1918. Components of the application need to be able to access other components of the application within the application's VPC by using the same host names as the host names that are used over the public internet. The network engineer also needs to accommodate future DNS changes, such as the introduction of new host names or the retirement of DNS entries. Which combination of steps will meet these requirements? (Choose three.)

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

Answer: BCD

NEW QUESTION 73

A retail company is running its service on AWS. The company's architecture includes Application Load Balancers (ALBs) in public subnets. The ALB target groups are configured to send traffic to backend Amazon EC2 instances in private subnets. These backend EC2 instances can call externally hosted services over the internet by using a NAT gateway. The company has noticed in its billing that NAT gateway usage has increased significantly. A network engineer needs to find out the source of this increased usage. Which options can the network engineer use to investigate the traffic through the NAT gateway? (Choose two.)

- A. Enable VPC flow logs on the NAT gateway's elastic network interfac
- B. Publish the logs to a log group in Amazon CloudWatch Log
- C. Use CloudWatch Logs Insights to query and analyze the logs.
- D. Enable NAT gateway access log
- E. Publish the logs to a log group in Amazon CloudWatch Log
- F. Use CloudWatch Logs Insights to query and analyze the logs.
- G. Configure Traffic Mirroring on the NAT gateway's elastic network interfac
- H. Send the traffic to an additional EC2 instanc
- I. Use tools such as tcpdump and Wireshark to query and analyze the mirrored traffic.
- J. Enable VPC flow logs on the NAT gateway's elastic network interfac
- K. Publish the logs to an Amazon S3 bucke
- L. Create a custom table for the S3 bucket in Amazon Athena to describe the log structur
- M. Use Athena to query and analyze the logs.
- N. Enable NAT gateway access log
- O. Publish the logs to an Amazon S3 bucke
- P. Create a custom table for the S3 bucket in Amazon Athena to describe the log structur
- Q. Use Athena to query and analyze the logs.

Answer: AD

Explanation:

To investigate the increased usage of a NAT gateway in a VPC architecture with ALBs and backend EC2 instances, a network engineer can use the following options:

➤ Enable VPC flow logs on the NAT gateway's elastic network interface and publish the logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query and analyze the logs. (Option A)

➤ Enable VPC flow logs on the NAT gateway's elastic network interface and publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure and use Athena to query and analyze the logs. (Option D)

These options allow for detailed analysis of traffic through the NAT gateway to identify the source of increased usage.

NEW QUESTION 74

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 77

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

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html>

To view all categories of instance metadata from within a running instance, use the following URI.

<http://169.254.169.254/latest/meta-data/>

NEW QUESTION 79

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

ANS-C01 Practice Exam Features:

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

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