



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

Answer: B

NEW QUESTION 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 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 endpoint
- 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 endpoint
- 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 endpoint
- 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 endpoint
- 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 5

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

Answer: A

NEW QUESTION 6

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 the anticipated 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 resolver
- 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 7

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 8

A real estate company is building an internal application so that real estate agents can upload photos and videos of various properties. The application will store these photos and videos in an Amazon S3 bucket as objects and will use Amazon DynamoDB to store corresponding metadata. The S3 bucket will be configured to publish all PUT events for new object uploads to an Amazon Simple Queue Service (Amazon SQS) queue.

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

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

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

Answer: C

NEW QUESTION 9

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 10

An insurance company is planning the migration of workloads from its on-premises data center to the AWS Cloud. The company requires end-to-end domain name resolution. Bi-directional DNS resolution between AWS and the existing on-premises environments must be established. The workloads will be migrated into multiple VPCs. The workloads also have dependencies on each other, and not all the workloads will be migrated at the same time. Which solution meets these requirements?

- A. Configure a private hosted zone for each application VPC, and create the requisite record
- B. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC
- C. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver
- D. Associate the application VPC private hosted zones with the egress VPC, and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager
- E. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints.
- F. Configure a public hosted zone for each application VPC, and create the requisite record
- G. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC
- H. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver
- I. Associate the application VPC private hosted zones with the egress VPC
- J. and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager
- K. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints.
- L. Configure a private hosted zone for each application VPC, and create the requisite record
- M. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC
- N. Associate the application VPC private hosted zones with the egress VPC and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager

Answer: A

Explanation:

Creating a private hosted zone for each application VPC and creating the requisite records would enable end-to-end domain name resolution for the resources. Creating a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC would enable bi-directional DNS resolution between AWS and the existing on-premises environments. Defining Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver would enable DNS queries from AWS resources to on-premises resources. Associating the application VPC private hosted zones with the egress VPC and sharing the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager would enable DNS queries among different VPCs and accounts. Configuring the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints would enable DNS queries from on-premises resources to AWS resources.

NEW QUESTION 10

A network engineer needs to set up an Amazon EC2 Auto Scaling group to run a Linux-based network appliance in a highly available architecture. The network engineer is configuring the new launch template for the Auto Scaling group. In addition to the primary network interface the network appliance requires a second network interface that will be used exclusively by the application to exchange traffic with hosts over the internet. The company has set up a Bring Your Own IP (BYOIP) pool that includes an Elastic IP address that should be used as the public IP address for the second network interface. How can the network engineer implement the required architecture?

- A. Configure the two network interfaces in the launch template
- B. Define the primary network interface to be created in one of the private subnets
- C. For the second network interface, select one of the public subnets
- D. Choose the BYOIP pool ID as the source of public IP addresses.
- E. Configure the primary network interface in a private subnet in the launch template
- F. Use the user data option to run a cloud-init script after boot to attach the second network interface from a subnet with auto-assign public IP addressing enabled.
- G. Create an AWS Lambda function to run as a lifecycle hook of the Auto Scaling group when an instance is launched
- H. In the Lambda function, assign a network interface to an AWS Global Accelerator endpoint.
- I. During creation of the Auto Scaling group, select subnets for the primary network interface
- J. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

Answer: D

Explanation:

During creation of the Auto Scaling group, select subnets for the primary network interface. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool. This solution meets all of the requirements stated in the question. The primary network interface can be configured in a private subnet during creation of the Auto Scaling group. The user data option can be used to run a cloud-init script that will allocate a second network interface and associate an Elastic IP address from the BYOIP pool with it.

NEW QUESTION 14

Your organization has a newly installed 1-Gbps AWS Direct Connect connection. You order the cross-connect from the Direct Connect location provider to the port on your router in the same facility. To enable the use of your first virtual interface, your router must be configured appropriately. What are the minimum requirements for your router?

- A. 1-Gbps Multi Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
- B. 1-Gbps Single Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
- C. IPsec Parameters, Pre-Shared key, Peer IP Address, BGP Session with MD5
- D. BGP Session with MD5, 802.1Q VLAN, Route-Map, Prefix List, IPsec encrypted GRE Tunnel

Answer: B

NEW QUESTION 15

A customer has set up multiple VPCs for Dev, Test, Prod, and Management. You need to set up AWS Direct Connect to enable data flow from on-premises to each VPC. The customer has monitoring software running in the Management VPC that collects metrics from the instances in all the other VPCs. Due to budget requirements, data transfer charges should be kept at minimum.

Which design should be recommended?

- A. Create a total of four private VIFs, one for each VPC owned by the customer, and route traffic between VPCs using the Direct Connect link.
- B. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs.
- C. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs, enable source/destination NAT in the Management VPC.
- D. Create a total of four private VIFs, and enable VPC peering between all VPCs.

Answer: D

Explanation:

- creating VPC peering is free of charge - traffic costs ~0.01€/GB for VPC peering (IN + OUT) and ~0.02€/GB for direct connect (OUT only). As the communication involved in monitoring will never have IN == OUT, then $0.01 * (IN + OUT)$ will always be lower than $0.02 * OUT$, ergo VPC peering will be cheaper

NEW QUESTION 19

A company has two AWS accounts one for Production and one for Connectivity. A network engineer needs to connect the Production account VPC to a transit gateway in the Connectivity account. The feature to auto accept shared attachments is not enabled on the transit gateway. Which set of steps should the network engineer follow in each AWS account to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 23

A company is building its website on AWS in a single VPC. The VPC has public subnets and private subnets in two Availability Zones. The website has static content such as images. The company is using Amazon S3 to store the content. The company has deployed a fleet of Amazon EC2 instances as web servers in a private subnet. The EC2 instances are in an Auto Scaling group behind an Application Load Balancer. The EC2 instances will serve traffic, and they must pull content from an S3 bucket to render the webpages. The company is using AWS Direct Connect with a public VIF for on-premises connectivity to the S3 bucket.

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

Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 26

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

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

What should a network engineer do to reduce this latency?

- A. Create a new Site-to-Site VPN connection
- B. Set the transit gateway as the target gateway

- C. Enable acceleration on the new Site-to-Site VPN 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 28

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 29

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 31

A company has deployed a software-defined WAN (SD-WAN) solution to interconnect all of its offices. The company is migrating workloads to AWS and needs to extend its SD-WAN solution to support connectivity to these workloads.

A network engineer plans to deploy AWS Transit Gateway Connect and two SD-WAN virtual appliances to provide this connectivity. According to company policies, only a single SD-WAN virtual appliance can handle traffic from AWS workloads at a given time.

How should the network engineer configure routing to meet these requirements?

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

Answer: A

NEW QUESTION 33

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

Which change should a network engineer implement to meet these requirements?

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

Answer: B

NEW QUESTION 38

A company is deploying an application. The application is implemented in a series of containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use the Fargate launch type for its tasks. The containers will run workloads that require connectivity initiated over an SSL connection. Traffic must be able to flow to the application from other AWS accounts over private connectivity. The application must scale in a manageable way as more consumers use the application.

Which solution will meet these requirements?

- A. Choose a Gateway Load Balancer (GLB) as the type of load balancer for the ECS service
- B. Create a lifecycle hook to add new tasks to the target group from Amazon ECS as required to handle scaling
- C. Specify the GLB in the service definition
- D. Create a VPC peer for external AWS account
- E. Update the route tables so that the AWS accounts can reach the GLB.
- F. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service
- G. Create path-based routing rules to allow the application to target the containers that are registered in the target group
- H. Specify the ALB in the service definition
- I. Create a VPC endpoint service for the ALB. Share the VPC endpoint service with other AWS accounts.
- J. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service
- K. Create path-based routing rules to allow the application to target the containers that are registered in the target group
- L. Specify the ALB in the service definition
- M. Create a VPC peer for the external AWS account
- N. Update the route tables so that the AWS accounts can reach the ALB.
- O. Choose a Network Load Balancer (NLB) as the type of load balancer for the ECS service
- P. Specify the NLB in the service definition
- Q. Create a VPC endpoint service for the NLB
- R. Share the VPC endpoint service with other AWS accounts.

Answer: D

NEW QUESTION 39

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

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

Which solution will meet these requirements?

- A. Deploy the EC2 instances in the public subnet
- B. Create an S3 interface endpoint in the VPC
- C. Modify the application configuration to use the S3 endpoint-specific DNS hostname.
- D. Deploy the EC2 instances in the private subnet
- E. Create a NAT gateway in the VPC
- F. Create default routes in the private subnets to the NAT gateway
- G. Connect to Amazon S3 by using the NAT gateway.
- H. Deploy the EC2 instances in the private subnet

- I. Create an S3 gateway endpoint in the VPC
- J. Specify the route table of the private subnets during endpoint creation to create routes to Amazon S3.
- K. Deploy the EC2 instances in the private subnet
- L. Create an S3 interface endpoint in the VPC
- M. 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 44

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 VPC
- 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.com
- E. Configure the private hosted zone to allow zone transfers with every VPC.
- F. Create one Amazon Route 53 private hosted zone for example.com
- G. Create a single resource record for aws.example.com in the private hosted zone
- H. Apply a multivalued answer routing policy to the records
- I. Add all VPC resources as separate values in the routing policy.
- J. Create one Amazon Route 53 private hosted zone for aws.example.com
- 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 46

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.monitoring
- 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 49

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

- A. Edit the existing Site-to-Site VPN connection by enabling acceleration
- B. Stop and start the VPN service on the customer gateway for the new setting to take effect.
- C. Configure a transit gateway in the same AWS Region as the existing virtual private gateway
- D. Create a new accelerated Site-to-Site VPN connection
- E. Connect the new connection to the transit gateway by using a VPN attachment
- F. Update the customer gateway device to use the new Site-to-Site VPN connection
- G. Delete the existing Site-to-Site VPN connection
- H. Create a new accelerated Site-to-Site VPN connection
- I. Connect the new Site-to-Site VPN connection to the existing virtual private gateway
- J. Update the customer gateway device to use the new Site-to-Site VPN connection
- 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 Cloud
- M. Update the customer gateway device to use the new Direct Connect connection
- N. Delete the existing Site-to-Site VPN connection.

Answer: B

NEW QUESTION 53

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

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

Answer: D

Explanation:

Configuring an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration would enable evaluation of the compliance status of the security groups based on predefined or custom 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 58

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 Region
- B. Create and attach private VIFs to a single Direct Connect gateway
- 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 center
- F. Share the transit gateway with each account by using AWS Resource Access Manager (AWS RAM). Attach the transit gateway to each VPC
- 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 center
- I. Share the transit gateways with each account by using AWS Resource Access Manager (AWS RAM). In each Region, attach the transit gateway to each VPC
- J. Peer all the VPCs in each Region to a new VPC in each Region that will function as a centralized transit VPC
- 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 61

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

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

What should the network engineer do to meet these requirements?

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

Answer: B

Explanation:

"If you advertise more than 100 routes each for IPv4 and IPv6 over the BGP session, the BGP session will go into an idle state with the BGP session DOWN."
<https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html>

NEW QUESTION 63

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

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

The network engineer must meet all the requirements for the transit gateway by designing a solution that needs the least number of transit gateway route tables. Which combination of actions should the network engineer perform to accomplish this goal?(Choose two.)

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

Answer: BD

NEW QUESTION 66

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 71

An international company provides early warning about tsunamis. The company plans to use IoT devices to monitor sea waves around the world. The data that is collected by the IoT devices must reach the company's infrastructure on AWS as quickly as possible. The company is using three operation centers around the world. Each operation center is connected to AWS through Its own AWS Direct Connect connection. Each operation center is connected to the internet through at least two upstream internet service providers.

The company has its own provider-independent (PI) address space. The IoT devices use TCP protocols for reliable transmission of the data they collect. The IoT devices have both landline and mobile internet connectivity. The infrastructure and the solution will be deployed in multiple AWS Regions. The company will use Amazon Route 53 for DNS services.

A network engineer needs to design connectivity between the IoT devices and the services that run in the AWS Cloud.

Which solution will meet these requirements with the HIGHEST availability?

- A. Set up an Amazon CloudFront distribution with origin failove
- B. Create an origin group for each Region where the solution is deployed.
- C. Set up Route 53 latency-based routin
- D. Add latency alias record
- E. For the latency alias records, set the value of Evaluate Target Health to Yes.
- F. Set up an accelerator in AWS Global Accelerato
- G. Configure Regional endpoint groups andhealth checks.
- H. Set up Bring Your Own IP (BYOIP) adresse
- I. Use the same PI addresses for each Region where the solution is deployed.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/iot/automate-global-device-provisioning-with-aws-iot-core-and-amazon-route-53>

NEW QUESTION 76

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 domain
- E. Use host-based routing rules to forward the traffic to the correct target group for each domain
- F. Include the X-Forwarded-For request header with traffic to the targets.
- G. Deploy a Network Load Balancer with a TLS listener
- H. Use path-based routing rules to forward the traffic to the correct target group
- I. Configure client IP address preservation for traffic to the targets.
- J. Deploy a Network Load Balancer with a TLS listener for each domain
- K. Use host-based routing rules to forward the traffic to the correct target group for each domain
- 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 80

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 81

A company operates its IT services through a multi-site hybrid infrastructure. The company deploys resources on AWS in the us-east-1 Region and in the eu-west-2 Region. The company also deploys resources in its own data centers that are located in the United States (US) and in the United Kingdom (UK). In both AWS Regions, the company uses a transit gateway to connect 15 VPCs to each other. The company has created a transit gateway peering connection between the two transit gateways. The VPC CIDR blocks do not overlap with each other or with IP addresses used within the data centers. The VPC CIDR prefixes can also be aggregated either on a Regional level or for the company's entire AWS environment.

The data centers are connected to each other by a private WAN connection. IP routing information is exchanged dynamically through Interior BGP (iBGP) sessions. The data centers maintain connectivity to AWS through one AWS Direct Connect connection in the US and one Direct Connect connection in the UK. Each Direct Connect connection is terminated on a Direct Connect gateway and is associated with a local transit gateway through a transit VIF.

Traffic follows the shortest geographical path from source to destination. For example, packets from the UK data center that are targeted to resources in eu-west-2 travel across the local Direct Connect connection. In cases of cross-Region data transfers, such as from the UK data center to VPCs in us-east-1, the private WAN connection must be used to minimize costs on AWS. A network engineer has configured each transit gateway association on the Direct Connect gateway to advertise VPC-specific CIDR IP prefixes only from the local Region. The routes toward the other Region must be learned through BGP from the routers in the other data center in the original, non-aggregated form.

The company recently experienced a problem with cross-Region data transfers because of issues with its private WAN connection. The network engineer needs to modify the routing setup to prevent similar interruptions in the future. The solution cannot modify the original traffic routing goal when the network is operating normally.

Which modifications will meet these requirements? (Choose two.)

- A. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connection
- B. Add the company's entire AWS environment aggregate route to the list of subnets advertised through the local Direct Connect connection.
- C. Add the CIDR prefixes from the other Region VPCs and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection
- D. Configure data center routers to make routing decisions based on the BGP communities received.
- E. Add the aggregate IP prefix for the other Region and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection.
- F. Add the aggregate IP prefix for the company's entire AWS environment and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection.
- G. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connection
- H. Add both Regional aggregate IP prefixes to the list of subnets advertised through the Direct Connect connection on both sides of the network
- I. Configure data center routers to make routing decisions based on the BGP communities received.

Answer: AD

NEW QUESTION 83

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 84

A Network Engineer is provisioning a subnet for a load balancer that will sit in front of a fleet of application servers in a private subnet. There is limited IP space left in the VPC CIDR. The application has few users now but is expected to grow quickly to millions of users.

What design will use the LEAST amount of IP space, while allowing for this growth?

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

Answer: C

NEW QUESTION 88

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 92

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 destinatio
- 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 93

A bank built a new version of its banking application in AWS using containers that connect 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 application
- 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 application
- 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 application
- 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 94

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 95

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

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

Answer: A

Explanation:

Setting up an AWS Site-to-Site VPN connection between on premises and AWS would enable a secure and encrypted connection over the public 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 98

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 rules. 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 ranges. Updating security groups to use the prefix list instead of the partner IP address range would enable simplification of security group rules. 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 list.

NEW QUESTION 100

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

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

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

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

Answer: BCE

Explanation:

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

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

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

NEW QUESTION 104

A security team is performing an audit of a company's AWS deployment. The security team is concerned that two applications might be accessing resources that should be blocked by network ACLs and security groups. The applications are deployed across two Amazon Elastic Kubernetes Service (Amazon EKS) clusters that use the Amazon VPC Container Network Interface (CNI) plugin for Kubernetes. The clusters are in separate subnets within the same VPC and have a Cluster Autoscaler configured.

The security team needs to determine which POD IP addresses are communicating with which services throughout the VPC. The security team wants to limit the number of flow logs and wants to examine the traffic from only the two applications.

Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

NEW QUESTION 108

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

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

Answer: BCD

NEW QUESTION 109

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 110

A company is running multiple workloads on Amazon EC2 instances in public subnets. In a recent incident, an attacker exploited an application vulnerability on one of the EC2 instances to gain access to the instance. The company fixed the application and launched a replacement EC2 instance that contains the updated application.

The attacker used the compromised application to spread malware over the internet. The company became aware of the compromise through a notification from AWS. The company needs the ability to identify when an application that is deployed on an EC2 instance is spreading malware.

Which solution will meet this requirement with the LEAST operational effort?

- A. Use Amazon GuardDuty to analyze traffic patterns by inspecting DNS requests and VPC flow logs.
- B. Use Amazon GuardDuty to deploy AWS managed decoy systems that are equipped with the most recent malware signatures.
- C. Set up a Gateway Load Balance
- D. Run an intrusion detection system (IDS) appliance from AWS Marketplace on Amazon EC2 for traffic inspection.
- E. Configure Amazon Inspector to perform deep packet inspection of outgoing traffic.

Answer: A

Explanation:

This solution involves using Amazon GuardDuty to monitor network traffic and analyze DNS requests and VPC flow logs for suspicious activity. This will allow the company to identify when an application is spreading malware by monitoring the network traffic patterns associated with the instance. GuardDuty is a fully managed threat detection service that continuously monitors for malicious activity and unauthorized behavior in your AWS accounts and workloads. It requires minimal setup and configuration and can be integrated with other AWS services for automated remediation. This solution requires the least operational effort compared to the other options.

NEW QUESTION 113

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 114

A network engineer is designing the architecture for a healthcare company's workload that is moving to the AWS Cloud. All data to and from the on-premises environment must be encrypted in transit. All traffic also must be inspected in the cloud before the traffic is allowed to leave the cloud and travel to the on-premises environment or to the internet.

The company will expose components of the workload to the internet so that patients can reserve appointments. The architecture must secure these components and protect them against DDoS attacks. The architecture also must provide protection against financial liability for services that scale out during a DDoS event. Which combination of steps should the network engineer take to meet all these requirements for the workload? (Choose three.)

- A. Use Traffic Mirroring to copy all traffic to a fleet of traffic capture appliances.
- B. Set up AWS WAF on all network components.
- C. Configure an AWS Lambda function to create Deny rules in security groups to block malicious IP addresses.
- D. Use AWS Direct Connect with MACsec support for connectivity to the cloud.
- E. Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection.
- F. Configure AWS Shield Advanced and ensure that it is configured on all public assets.

Answer: DEF

Explanation:

To meet the requirements for the healthcare company's workload that is moving to the AWS Cloud, the network engineer should take the following steps:

- Use AWS Direct Connect with MACsec support for connectivity to the cloud to ensure that all data to and from the on-premises environment is encrypted in transit (Option D).
- Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection to inspect all traffic in the cloud before it is allowed to leave (Option E).
- Configure AWS Shield Advanced and ensure that it is configured on all public assets to secure components exposed to the internet against DDoS attacks and provide protection against financial liability for services that scale out during a DDoS event (Option F).

These steps will help ensure that all data is encrypted in transit, all traffic is inspected before leaving the cloud, and components exposed to the internet are secured against DDoS attacks.

NEW QUESTION 119

.....

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](#)