



Exam Code: ANS-C01

Exam Name: AWS Certified Advanced Networking - Specialty (ANS-C01)

Website: www.VCEplus.io - www.VCEup.com

VCEup

Question No: 1

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

Which solution will meet these requirements?

- A. Install the AWS Load Balancer Controller for Kubernetes. Using that controller, configure a Network Load Balancer with a TCP listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
- B. Install the AWS Load Balancer Controller for Kubernetes. Using that controller, configure an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
- C. Create a target group. Add the EKS managed node group's Auto Scaling group as a target. Create an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the target group.
- D. Create a target group. Add the EKS managed node group's Auto Scaling group as a target. Create a Network Load Balancer with a TLS listener on port 443 to forward traffic to the target group.

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-targetgroups.html#target-group-protocol-version> <https://docs.aws.amazon.com/prescriptiveguidance/latest/patterns/deploy-a-grpc-based-application-on-an-amazon-eks-cluster-and-access-it-with-an-application-load-balancer.html>

Question No: 2

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 listener. Use path-based routing rules to forward the traffic to the correct target group. Include the X-Forwarded-For request header with traffic to the targets.
- B. Deploy an Application Load Balancer with an HTTPS listener for each domain. Use host-based routing rules to forward the traffic to the correct target group for each domain. Include the X-Forwarded-For request header with traffic to the targets.
- C. Deploy a Network Load Balancer with a TLS listener. Use path-based routing rules to forward the traffic to the correct target group. Configure client IP address preservation for traffic to the targets.
- D. Deploy a Network Load Balancer with a TLS listener for each domain. Use host-based routing rules to forward the traffic to the correct target group for each domain. 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.

Question No: 3

A company has developed an application on AWS that will track inventory levels of vending machines and initiate the restocking process automatically. The company plans to integrate this application with vending machines and deploy the vending machines in several markets around the world. The application resides in a VPC in the us-east-1 Region. The application consists of an Amazon Elastic Container Service (Amazon ECS) cluster behind an Application Load Balancer (ALB). The communication from the vending machines to the application happens over HTTPS.

The company is planning to use an AWS Global Accelerator accelerator and configure static IP addresses of the accelerator in the vending machines for application endpoint access. The application must be accessible only through the accelerator and not through a direct connection over the internet to the ALB endpoint.

Which solution will meet these requirements?

- A. Configure the ALB in a private subnet of the VPC. Attach an internet gateway without adding routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- B. Configure the ALB in a private subnet of the VPC. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- C. Configure the ALB in a public subnet of the VPC. Attach an internet gateway. Add routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.
- D. Configure the ALB in a private subnet of the VPC. Attach an internet gateway. Add routes in the subnet route tables to point to the internet gateway. Configure the accelerator with endpoint groups that include the ALB endpoint. Configure the ALB's security group to only allow inbound traffic from 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 accelerator (and the last line of the extract) - <https://docs.aws.amazon.com/global-accelerator/latest/dg/secure-vpcconnections.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."

Question No: 4

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

Question No: 5

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 observed. Create a new 10 Gbps dedicated connection. Shift traffic from the existing dedicated connection to the new dedicated connection.
- B. 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. Upgrade the bandwidth of the existing dedicated connection to 10 Gbps.
- C. 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. Upgrade the existing dedicated connection to a 5 Gbps hosted connection.
- D. 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 connection. 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).

Question No: 6

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

A recent design meeting revealed that the customers have IP address overlap with the provider's AWS deployment. The customers have stated that they will not share their internal IP addresses and that they do not want to connect to the provider's SaaS service over the internet.

Which combination of steps is part of a solution that meets these requirements? (Choose two.)

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

Answer: AB

Explanation:

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

Question No: 7

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.

Question No: 8

A retail company is running its service on AWS. The company's architecture includes Application Load Balancers (ALBs) in public subnets. The ALB target groups are configured to send traffic to backend Amazon EC2 instances in private subnets. These backend EC2 instances can call externally hosted services over the internet by using a NAT gateway.

The company has noticed in its billing that NAT gateway usage has increased significantly. A network engineer needs to find out the source of this increased usage.

Which options can the network engineer use to investigate the traffic through the NAT gateway?

(Choose two.)

- A. Enable VPC flow logs on the NAT gateway's elastic network interface. Publish the logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query and analyze the logs.
- B. Enable NAT gateway access logs. Publish the logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query and analyze the logs.
- C. Configure Traffic Mirroring on the NAT gateway's elastic network interface. Send the traffic to an additional EC2 instance. Use tools such as tcpdump and Wireshark to query and analyze the mirrored traffic.
- D. Enable VPC flow logs on the NAT gateway's elastic network interface. Publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure. Use Athena to query and analyze the logs.
- E. Enable NAT gateway access logs. Publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure. Use Athena to query and analyze the logs.

Answer: AD

Explanation:

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

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

(Option A)

Enable VPC flow logs on the NAT gateway's elastic network interface and publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure and use Athena to query and analyze the logs. (Option D) These options allow for detailed analysis of traffic through the NAT gateway to identify the source of increased usage.

Question No: 9

A banking company is successfully operating its public mobile banking stack on AWS. The mobile banking stack is deployed in a VPC that includes private subnets and public subnets. The company is using IPv4 networking and has not deployed or supported IPv6 in the environment. The company has decided to adopt a third-party service provider's API and must integrate the API with the existing environment. The service provider's API requires the use of IPv6.

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

Question No: 10

A company has deployed an AWS Network Firewall firewall into a VPC. A network engineer needs to implement a solution to deliver Network Firewall flow logs to the company's Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster in the shortest possible time.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket. Create an AWS Lambda function to load logs into the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster. Enable Amazon Simple Notification Service (Amazon SNS) notifications on the S3 bucket to invoke the Lambda function. Configure flow logs for the firewall. Set the S3 bucket as the destination.
- B. Create an Amazon Kinesis Data Firehose delivery stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination. Configure flow logs for the firewall. Set the Kinesis Data Firehose delivery stream as the destination for the Network Firewall flow logs.
- C. Configure flow logs for the firewall. Set the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination for the Network Firewall flow logs.
- D. Create an Amazon Kinesis data stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination. Configure flow logs for the firewall. Set the Kinesis data stream as the destination for the Network Firewall flow logs.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-analyze-aws-networkfirewall-logs-using-amazon-opensearch-service-part-1/>

Question No: 11

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 for efs.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 VPC. Update all the VPC DHCP options sets to use AmazonProvidedDNS for name resolution.
- C. Create an Amazon Route 53 Resolver inbound endpoint in the central VPC. Update all the VPC DHCP options sets to use the Route 53 Resolver inbound endpoint in the central VPC for name resolution.
- D. Create an Amazon Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS servers. Share the rule with the organization by using AWS Resource Access Manager (AWS RAM). Associate the rule with all the VPCs.
- E. 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 deployed. Create an A record for fs-33444567d.efs.us-east-1.amazonaws.com in the private hosted zone. 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.

Question No: 12

An ecommerce company is hosting a web application on Amazon EC2 instances to handle continuously changing customer demand. The EC2 instances are part of an Auto Scaling group. The company wants to implement a solution to distribute traffic from customers to the EC2 instances.

The company must encrypt all traffic at all stages between the customers and the application servers.

No decryption at intermediate points is allowed.

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

Question No: 13

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. 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>

Question No: 14

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 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.

- C. Create an Amazon Route 53 Resolver rule to forward any queries made to onprem.example.internal to the on-premises DNS servers.
- D. Create an Amazon Route 53 private hosted zone named aws.example.internal in the shared AWS account to resolve queries for this domain.
- E. Launch two Amazon EC2 instances in the shared AWS account. Install BIND on each instance.

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 account. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS servers.

Set the forwarding IP addresses to the IP addresses of the BIND servers.

- F. 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 costeffectiveness 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.

Question No: 15

A company has multiple AWS accounts. Each account contains one or more VPCs. A new security guideline requires the inspection of all traffic between VPCs.

The company has deployed a transit gateway that provides connectivity between all VPCs. The company also has deployed a shared services VPC with Amazon EC2 instances that include IDS services for stateful inspection. The EC2 instances are deployed across three Availability Zones. The company has set up VPC associations and routing on the transit gateway. The company has migrated a few test VPCs to the new solution for traffic inspection.

Soon after the configuration of routing, the company receives reports of intermittent connections for traffic that crosses Availability Zones.

What should a network engineer do to resolve this issue?

A. Modify the transit gateway VPC attachment on the shared services VPC by enabling cross- Availability Zone load balancing.

B. Modify the transit gateway VPC attachment on the shared services VPC by enabling appliance mode support.

C. Modify the transit gateway by selecting VPN equal-cost multi-path (ECMP) routing support.

D. Modify the transit gateway by selecting multicast support.

Answer: B

Explanation:

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

Question No: 16

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. Associate the new security group with the endpoint network interfaces.

E. 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.

F. Associate the VPC endpoint or endpoints with route tables that the private subnets use.

Answer: B, D, F

Explanation:

Question No: 17

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

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

The company will use Amazon Route 53 for DNS services.

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

Which solution will meet these requirements with the HIGHEST availability?

A. Set up an Amazon CloudFront distribution with origin failover. Create an origin group for each Region where the solution is deployed.

B. Set up Route 53 latency-based routing. Add latency alias records. For the latency alias records, set the value of Evaluate Target Health to Yes.

C. Set up an accelerator in AWS Global Accelerator. Configure Regional endpoint groups and health checks.

D. Set up Bring Your Own IP (BYOIP) addresses. 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-andamazon-route-53/>

Question No: 18

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

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

A. Configure a public VIF on the Direct Connect connection. Configure an AWS Site-to-Site VPN connection to the transit gateway as a VPN attachment.

B. Configure a transit VIF on the Direct Connect connection. Configure an IPsec VPN connection to an EC2 instance that is running third-party VPN software.

C. Configure MACsec for the Direct Connect connection. Configure a transit VIF to a Direct Connect gateway that is associated with the transit gateway.

D. Configure a public VIF on the Direct Connect connection. Configure two AWS Site-to-Site VPN connections to the transit gateway. Enable equal-cost multi-path (ECMP) routing.

Answer: C

Explanation:

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

Question No: 19

A network engineer must develop an AWS CloudFormation template that can create a virtual private gateway, a customer gateway, a VPN connection, and static routes in a route table. During testing of the template, the network engineer notes that the CloudFormation template has encountered an error and is rolling back.

What should the network engineer do to resolve the error?

- A. Change the order of resource creation in the CloudFormation template.
- B. Add the DependsOn attribute to the resource declaration for the virtual private gateway. Specify the route table entry resource.
- C. Add a wait condition in the template to wait for the creation of the virtual private gateway.
- D. Add the DependsOn attribute to the resource declaration for the route table entry. Specify the virtual private gateway resource.

Answer: D

Explanation:

Question No: 20

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 useast-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. Add the company's entire AWS environment aggregate route to the list of subnets advertised through the local Direct Connect connection.
- B. 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. Configure data center routers to make routing decisions based on the BGP communities received.
- C. 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.
- D. 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.
- E. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connection. Add both Regional aggregate IP prefixes to the list of subnets advertised through the Direct Connect connection on both sides of the network. Configure data center routers to make routing decisions based on the BGP communities received.

Answer: AD

Explanation:

Question No: 21

A company's network engineer needs to design a new solution to help troubleshoot and detect network anomalies. The network engineer has configured Traffic Mirroring. However, the mirrored traffic is overwhelming the Amazon EC2 instance that is the traffic mirror target. The EC2 instance hosts tools that the company's security team uses to analyze the traffic. The network engineer needs to design a highly available solution that can scale to meet the demand of the mirrored traffic.

Which solution will meet these requirements?

- A. Deploy a Network Load Balancer (NLB) as the traffic mirror target. Behind the NLB, deploy a fleet of EC2 instances in an Auto Scaling group. Use Traffic Mirroring as necessary.
- B. Deploy an Application Load Balancer (ALB) as the traffic mirror target. Behind the ALB, deploy a fleet of EC2 instances in an Auto Scaling group. Use Traffic Mirroring only during non-business hours.
- C. Deploy a Gateway Load Balancer (GLB) as the traffic mirror target. Behind the GLB, deploy a fleet of EC2 instances in an Auto Scaling group. Use Traffic Mirroring as necessary.
- D. Deploy an Application Load Balancer (ALB) with an HTTPS listener as the traffic mirror target.

Behind the ALB, deploy a fleet of EC2 instances in an Auto Scaling group. Use Traffic Mirroring only during active events or business hours.

Answer: A

Explanation:

Question No: 22

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.

Question No: 23

A government contractor is designing a multi-account environment with multiple VPCs for a customer. A network security policy requires all traffic between any two VPCs to be transparently inspected by a third-party appliance.

The customer wants a solution that features AWS Transit Gateway. The setup must be highly available across multiple Availability Zones, and the solution needs to support automated failover.

Furthermore, asymmetric routing is not supported by the inspection appliances.

Which combination of steps is part of a solution that meets these requirements? (Choose two.)

A. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VPC. Connect the inspection VPC to the transit gateway by using a VPC attachment. Create a target group, and register the appliances with the target group. Create a Network Load Balancer (NLB), and set it up to forward to the newly created target group. Configure a default route in the inspection VPCs transit gateway subnet toward the NLB.

B. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VPC. Connect the inspection VPC to the transit gateway by using a VPC attachment. Create a target group, and register the appliances with the target group. Create a Gateway Load Balancer, and set it up to forward to the newly created target group. Configure a default route in the inspection VPC's transit gateway subnet toward the Gateway Load Balancer endpoint.

C. Configure two route tables on the transit gateway. Associate one route table with all the attachments of the application VPCs. Associate the other route table with the inspection VPC's attachment. Propagate all VPC attachments into the inspection route table. Define a static default route in the application route table. Enable appliance mode on the attachment that connects the inspection VPC.

D. Configure two route tables on the transit gateway. Associate one route table with all the attachments of the application VPCs. Associate the other route table with the inspection VPCs attachment. Propagate all VPC attachments into the application route table. Define a static default route in the inspection route table. Enable appliance mode on the attachment that connects the inspection VPC.

E. Configure one route table on the transit gateway. Associate the route table with all the VPCs.

Propagate all VPC attachments into the route table. Define a static default route in the route table.

Answer: B, C

Explanation:

Question No: 24

A company has deployed Amazon EC2 instances in private subnets in a VPC. The EC2 instances must initiate any requests that leave the VPC, including requests to the company's on-premises data center over an AWS Direct Connect connection. No resources outside the VPC can be allowed to open communications directly to the EC2 instances.

The on-premises data center's customer gateway is configured with a stateful firewall device that filters for incoming and outgoing requests to and from multiple VPCs. In addition, the company wants to use a single IP match rule to allow all the communications from the EC2 instances to its data center from a single IP address.

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

A. Create a VPN connection over the Direct Connect connection by using the on-premises firewall.

Use the firewall to block all traffic from on premises to AWS. Allow a stateful connection from the EC2 instances to initiate the requests.

B. Configure the on-premises firewall to filter all requests from the on-premises network to the EC2 instances. Allow a stateful connection if the EC2 instances in the VPC initiate the traffic.

C. Deploy a NAT gateway into a private subnet in the VPC where the EC2 instances are deployed.

Specify the NAT gateway type as private. Configure the on-premises firewall to allow connections from the IP address that is assigned to the NAT gateway.

D. Deploy a NAT instance into a private subnet in the VPC where the EC2 instances are deployed.

Configure the on-premises firewall to allow connections from the IP address that is assigned to the NAT instance.

Answer: C

Explanation:

Question No: 25

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. Create and attach private VIFs to a single Direct Connect gateway. Attach the Direct Connect gateway to all the VPCs. Remove the existing VPN connections that are attached directly to the virtual private gateways.
- B. Create a single transit gateway with VPN connections from each data center. Share the transit gateway with each account by using AWS Resource Access Manager (AWS RAM). Attach the transit gateway to each VPC. Remove the existing VPN connections that are attached directly to the virtual private gateways.
- C. Create a transit gateway in each Region with multiple newly commissioned VPN connections from each data center. 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. Remove the existing VPN connections that are attached directly to the virtual private gateways.
- D. Peer all the VPCs in each Region to a new VPC in each Region that will function as a centralized transit VPC. Create new VPN connections from each data center to the transit VPCs. Terminate the original VPN connections that are attached to all the original VPCs. Retain the new VPN connection to the new transit VPC in each Region.

Answer: C

Explanation:

Question No: 26

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

Answer: D

Explanation:

Question No: 27

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 routes. Configure a CloudWatch alarm to send notifications when routes change.

- B. Onboard Transit Gateway Network Manager to Amazon CloudWatch Logs Insights. Use Amazon EventBridge (Amazon CloudWatch Events) to send notifications when routes change.
- C. Configure an AWS Lambda function to periodically check the routes on the Direct Connect gateway and to send notifications when routes change.
- D. Enable Amazon CloudWatch Logs on the transit VIFs to track the received routes. 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-cloudwatchevents.html>To receive notification each time a new route is advertised to AWS from on premises over DirectConnect, a network engineer should onboard Transit Gateway Network Manager to AmazonCloudWatch Logs Insights and use Amazon EventBridge (Amazon CloudWatch Events) to sendnotifications when routes change (Option B). This solution allows for real-time monitoring of routechanges and automatic notification when new routes are advertised.

Question No: 28

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 onpremises 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 connections. Reroute traffic through the new public VIF.
- B. 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.
- C. Deploy a new pair of 10 GB Direct Connect connections with MACsec. Configure MACsec on the edge routers. Reroute traffic to the new Direct Connect connections. Decommission the original Direct Connect connections
- D. Deploy a new pair of 10 GB Direct Connect connections with MACsec. Deploy a new public VIF on the new Direct Connect connections. Deploy two AWS Site-to-Site VPN connections on top of the new public VIF. Reroute traffic from the existing private VIF to the new Site-to-Site connections.

Decommission the original Direct Connect connections.

Answer: C

Explanation:

Question No: 29

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. Create an AWS Lambda function as the target of the rule. Configure the function to use the event information to update the private hosted zone.

F. Add the private IP addresses in the existing Route 53 public hosted zone.

Answer: B, C, D

Explanation:

Question No: 30

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. Create a lifecycle hook to add new tasks to the target group from Amazon ECS as required to handle scaling.

Specify the GLB in the service definition. Create a VPC peer for external AWS accounts. Update the route tables so that the AWS accounts can reach the GLB.

B. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service. Create path-based routing rules to allow the application to target the containers that are registered in the target group. Specify the ALB in the service definition. Create a VPC endpoint service for the ALB. Share the VPC endpoint service with other AWS accounts.

C. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service. Create path-based routing rules to allow the application to target the containers that are registered in the target group. Specify the ALB in the service definition. Create a VPC peer for the external AWS accounts. Update the route tables so that the AWS accounts can reach the ALB.

D. Choose a Network Load Balancer (NLB) as the type of load balancer for the ECS service. Specify the NLB in the service definition. Create a VPC endpoint service for the NLB. Share the VPC endpoint service with other AWS accounts.

Answer: D

Explanation:

Question No: 31

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

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

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

Which solution will meet these requirements?

A. Create a VPC peering connection between the web service VPC and the existing production VPC.

Add a routing rule to the appropriate route table to allow data to flow to 192.168.224.0/19 from the existing production environment and to flow to 192.168.128.0/17 from the web service environment.

Configure the relevant security groups and ACLs to allow the systems to communicate.

B. Ask the development team of the web service to redeploy the web service into the production VPC and integrate the systems there.

C. Create a VPC endpoint service. Associate the VPC endpoint service with the NLB for the web service. Create an interface VPC endpoint for the web service in the existing production VPC.

D. Create a transit gateway in the existing production environment. Create attachments to the production VPC and the web service VPC. Configure appropriate routing rules in the transit gateway and VPC route tables for 192.168.224.0/19 and 192.168.128.0/17. Configure the relevant security groups and ACLs to allow the systems to communicate.

Answer: C

Explanation:

Question No: 32

A network engineer needs to update a company's hybrid network to support IPv6 for the upcoming release of a new application. The application is hosted in a VPC in the AWS Cloud. The company's current AWS infrastructure includes VPCs that are connected by a transit gateway. The transit gateway is connected to the on-premises network by AWS Direct Connect and AWS Site-to-Site VPN.

The company's on-premises devices have been updated to support the new IPv6 requirements.

The company has enabled IPv6 for the existing VPC by assigning a new IPv6 CIDR block to the VPC and by assigning IPv6 to the subnets for dual-stack support. The company has launched new Amazon EC2 instances for the new application in the updated subnets.

When updating the hybrid network to support IPv6 the network engineer must avoid making any changes to the current infrastructure. The network engineer also must block direct access to the instances' new IPv6 addresses from the internet. However, the network engineer must allow outbound internet access from the instances.

What is the MOST operationally efficient solution that meets these requirements?

- A. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Create a new VPN connection that supports IPv6 connectivity. Add an egress-only internet gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices
- B. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Update the existing VPN connection to support IPv6 connectivity. Add an egress-only internet gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
- C. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Create a new VPN connection that supports IPv6 connectivity. Add an egress-only internet gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
- D. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address. Create a new VPN connection that supports IPv6 connectivity. Add a NAT gateway. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.

Answer: B

Explanation:

Question No: 33

A network engineer must provide additional safeguards to protect encrypted data at Application Load Balancers (ALBs) through the use of a unique random session key.

What should the network engineer do to meet this requirement?

- A. Change the ALB security policy to a policy that supports TLS 1.2 protocol only
- B. Use AWS Key Management Service (AWS KMS) to encrypt session keys
- C. Associate an AWS WAF web ACL with the ALBs. and create a security rule to enforce forward secrecy (FS)
- D. Change the ALB security policy to a policy that supports forward secrecy (FS)

Answer: D

Explanation:

Question No: 34

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

Answer: A

Explanation:

Question No: 35

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

Which solution will meet these requirements?

- A. Create a new VPC for the SD-WAN hub virtual appliance. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway. Configure BGP over the IPsec VPN connections
- B. Assign a new CIDR block to the transit gateway. Create a new VPC for the SD-WAN hub virtual appliance. Attach the new VPC to the transit gateway with a VPC attachment. Add a transit gateway Connect attachment. Create a Connect peer and specify the GRE and BGP parameters. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.
- C. Create a new VPC for the SD-WAN hub virtual appliance. Attach the new VPC to the transit gateway with a VPC attachment. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway. Configure BGP over the IPsec VPN connections.
- D. Assign a new CIDR block to the transit gateway. Create a new VPC for the SD-WAN hub virtual appliance. Attach the new VPC to the transit gateway with a VPC attachment. Add a transit gateway Connect attachment. Create a Connect peer and specify the VXLAN and BGP parameters. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.

Answer: D

Explanation:

Question No: 36

A company is deploying a new application on AWS. The application uses dynamic multicasting. The company has five VPCs that are all attached to a transit gateway Amazon EC2 instances in each VPC need to be able to register dynamically to receive a multicast transmission.

How should a network engineer configure the AWS resources to meet these requirements?

- A. Create a static source multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
- B. Create a static source multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.
- C. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
- D. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain. Register the multicast senders' network interface with the multicast domain. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.

Answer: C

Explanation:

Question No: 37

A company has deployed a critical application on a fleet of Amazon EC2 instances behind an Application Load Balancer. The application must always be reachable on port 443 from the public internet. The application recently had an outage that resulted from an incorrect change to the EC2 security group.

A network engineer needs to automate a way to verify the network connectivity between the public internet and the EC2 instances whenever a change is made to the security group. The solution also must notify the network engineer when the change affects the connection.

Which solution will meet these requirements?

- A. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture REJECT traffic on port 443. Publish the flow log records to a log group in Amazon CloudWatch Logs. Create a CloudWatch Logs metric filter for the log group for rejected traffic. Create an alarm to notify the network engineer.
- B. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture all traffic on port 443. Publish the flow log records to a log group in Amazon CloudWatch Logs. Create a CloudWatch Logs metric filter for the log group for all traffic. Create an alarm to notify the network engineer
- C. Create a VPC Reachability Analyzer path on port 443. Specify the security group as the source.

Specify the EC2 instances as the destination. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs.

D. Create a VPC Reachability Analyzer path on port 443. Specify the internet gateway of the VPC as the source. Specify the EC2 instances as the destination. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs.

Answer: C

Explanation:

Question No: 38

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 format. Create a filter to gather flow logs only from the EKS nodes. Include the srcaddr field and the dstaddr field in the flow logs.
- B. Create VPC flow logs in a custom format. Set the EKS nodes as the resource. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.
- C. Create VPC flow logs in a custom format. Set the application subnets as resources. Include the pktsrcaddr field and the pkt-dstaddr field in the flow logs.
- D. Create VPC flow logs in a custom format. 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

Explanation:

Question No: 39

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 VPC. Reconfigure half of the HPC cluster nodes to use these new DNS servers. Plan to scale out by adding additional EC2 instance-based DNS servers in the future as the HPC cluster size grows.
- B. Scale up the existing EC2 instances that the company is using as DNS servers. Change the instance size to the largest possible instance size to accommodate the current DNS load and the anticipated load in the future.
- C. Create Route 53 Resolver outbound endpoints. Create Route 53 Resolver rules to forward queries to on-premises DNS servers for on-premises hosted domain names. Reconfigure the HPC cluster nodes to use the default VPC resolver instead of the EC2 instance-based DNS servers. Terminate the EC2 instances.
- D. Create Route 53 Resolver inbound endpoints. Create rules on the on-premises DNS servers to forward queries to the default VPC resolver. Reconfigure the HPC cluster nodes to forward all DNS queries to the on-premises DNS servers. Terminate the EC2 instances.

Answer: C

Explanation:

Question No: 40

A company's network engineer is designing an active-passive connection to AWS from two on-premises data centers. The company has set up AWS Direct Connect connections between the on-premises data centers and AWS. From each location, the company is using a transit VIF that connects to a Direct Connect gateway that is associated with a transit gateway.

The network engineer must ensure that traffic from AWS to the data centers is routed first to the primary data center. The traffic should be routed to the failover data center only in the case of an outage.

Which solution will meet these requirements?

- A. Set the BGP community tag for all prefixes from the primary data center to 7224:7100. Set the BGP community tag for all prefixes from the failover data center to 7224:7300
- B. Set the BGP community tag for all prefixes from the primary data center to 7224:7300. Set the BGP community tag for all prefixes from the failover data center to 7224:7100
- C. Set the BGP community tag for all prefixes from the primary data center to 7224:9300. Set the BGP community tag for all prefixes from the failover data center to 7224:9100
- D. Set the BGP community tag for all prefixes from the primary data center to 7224:9100. Set the BGP community tag for all prefixes from the failover data center to 7224:9300

Answer: B

Explanation:

Question No: 41

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. Disable the Auto-assign Public IP option while launching the EC2 instances. Create an internet gateway. Attach the internet gateway to the VPC. In the public subnet's route table, add a default route that points to the internet gateway.
- B. Place the EC2 instances in a private subnet. Create a NAT gateway in a public subnet in the same Availability Zone. Create an internet gateway. Attach the internet gateway to the VPC. In the public subnet's route table, add a default route that points to the internet gateway.

C. Place the EC2 instances in a private subnet. Create an interface VPC endpoint for Amazon SQS.

Create gateway VPC endpoints for Amazon S3 and DynamoDB.

D. Place the EC2 instances in a private subnet. Create a gateway VPC endpoint for Amazon SQS.

Create interface VPC endpoints for Amazon S3 and DynamoDB.

Answer: C

Explanation:

Question No: 42

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

The company is opening a new office in Europe with a new on-premises data center in England. A Direct Connect connection will connect the new data center with some workloads that are running in a single VPC in the eu-west-2 Region. The company needs to connect the US data center and us-east-1 with the Europe data center and eu-west-2. A network engineer must establish full connectivity between the data centers and Regions with the lowest possible latency.

How should the network engineer design the network architecture to meet these requirements?

A. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VIF. Associate the transit gateway in us-east-1 with the same Direct Connect gateway.

Enable SiteLink for the transit VIF and the private VIF.

B. Connect the VPC in eu-west-2 to a new transit gateway. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VIF. Associate the transit gateway in us-east-1 with the same Direct Connect gateway. Enable SiteLink for both transit VIFs.

Peer the two transit gateways.

C. Connect the VPC in eu-west-2 to a new transit gateway. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VIF. Create a new Direct Connect gateway. Associate the transit gateway in us-east-1 with the new Direct Connect gateway.

Enable SiteLink for both transit VIFs. Peer the two transit gateways.

D. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VIF. Create a new Direct Connect gateway. Associate the transit gateway in us-east-1 with the new Direct Connect gateway. Enable SiteLink for the transit VIF and the private VIF.

Answer: C

Explanation:

Question No: 43

A network engineer has deployed an Amazon EC2 instance in a private subnet in a VPC. The VPC has no public subnet. The EC2 instance hosts application code that sends messages to an Amazon Simple Queue Service (Amazon SQS) queue. The subnet has the default network ACL with no modification applied. The EC2 instance has the default security group with no modification applied.

The SQS queue is not receiving messages.

Which of the following are possible causes of this problem? (Choose two.)

A. The EC2 instance is not attached to an IAM role that allows write operations to Amazon SQS.

B. The security group is blocking traffic to the IP address range used by Amazon SQS

C. There is no interface VPC endpoint configured for Amazon SQS

D. The network ACL is blocking return traffic from Amazon SQS

E. There is no route configured in the subnet route table for the IP address range used by Amazon SQS

Answer: C, E

Explanation:

Question No: 44

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 this centralized 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 KMS. Modify the interface endpoint by disabling the private DNS name. Create a private hosted zone in the shared services account with an alias record that points to the interface endpoint. Associate the private hosted zone with the spoke VPCs in each AWS account.

B. In the shared services account, create an interface endpoint for AWS KMS. Modify the interface endpoint by disabling the private DNS name. Create a private hosted zone in each spoke AWS account with an alias record that points to the interface endpoint. Associate each private hosted zone with the shared services AWS account.

C. In each spoke AWS account, create an interface endpoint for AWS KMS. Modify each interface endpoint by disabling the private DNS name. Create a private hosted zone in each spoke AWS account with an alias record that points to each interface endpoint. Associate each private hosted zone with the shared services AWS account.

D. In each spoke AWS account, create an interface endpoint for AWS KMS. Modify each interface endpoint by disabling the private DNS name. Create a private hosted zone in the shared services account with an alias record that points to each interface endpoint. Associate the private hosted zone with the spoke VPCs in each AWS account.

Answer: A

Explanation:

Question No: 45

A company has created three VPCs: a production VPC, a nonproduction VPC, and a shared services VPC. The production VPC and the nonproduction VPC must each have communication with the shared services VPC. There must be no communication between the production VPC and the nonproduction VPC. A transit gateway is deployed to facilitate communication between VPCs.

Which route table configurations on the transit gateway will meet these requirements?

A. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for only the shared services VPC. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.

B. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for each VPC. Create an additional route table with only the shared services VPC attachment associated with propagated routes from each VPC.

C. Configure a route table with all the VPC attachments associated with propagated routes for only the shared services VPC. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.

D. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes disabled. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.

Answer: A

Explanation:

Question No: 46

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. Stop and start the VPN service on the customer gateway for the new setting to take effect.
- B. Configure a transit gateway in the same AWS Region as the existing virtual private gateway. Create a new accelerated Site-to-Site VPN connection. Connect the new connection to the transit gateway by using a VPN attachment. Update the customer gateway device to use the new Site to Site VPN connection. Delete the existing Site-to-Site VPN connection
- C. Create a new accelerated Site-to-Site VPN connection. Connect the new Site-to-Site VPN connection to the existing virtual private gateway. Update the customer gateway device to use the new Site-to-Site VPN connection. Delete the existing Site-to-Site VPN connection.
- D. Create a new AWS Direct Connect connection with a private VIF between the on-premises data center and the AWS Cloud. Update the customer gateway device to use the new Direct Connect connection. Delete the existing Site-to-Site VPN connection.

Answer: B

Explanation:

Question No: 47

An Australian ecommerce company hosts all of its services in the AWS Cloud and wants to expand its customer base to the United States (US). The company is targeting the western US for the expansion.

The company's existing AWS architecture consists of four AWS accounts with multiple VPCs deployed in the ap-southeast-2 Region. All VPCs are attached to a transit gateway in ap-southeast-2. There are dedicated VPCs for each application service. The company also has VPCs for centralized security features such as proxies, firewalls, and logging.

The company plans to duplicate the infrastructure from ap-southeast-2 to the us-west-1 Region. A network engineer must establish connectivity between the various applications in the two Regions.

The solution must maximize bandwidth, minimize latency and minimize operational overhead.

Which solution will meet these requirements?

- A. Create VPN attachments between the two transit gateways. Configure the VPN attachments to use BGP routing between the two transit gateways.
- B. Peer the transit gateways in each Region. Configure routing between the two transit gateways for each Region's IP addresses.
- C. Create a VPN server in a VPC in each Region. Update the routing to point to the VPN servers for the IP addresses in alternate Regions.
- D. Attach the VPCs in us-west-1 to the transit gateway in ap-southeast-2.

Answer: B

Explanation:

Peering the transit gateways in each region would establish a private network connection between the two regions, allowing the company to route traffic between the VPCs in different regions without going over the public internet. This would help minimize latency and maximize bandwidth while reducing the operational overhead of managing multiple VPN connections.

Question No: 48

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 listeners. Use Bring Your Own IP (BYOIP) from the on-premises network with the NLB.
- B. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listeners. Create an AWS Global Accelerator accelerator in front of the NLB. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator.
- C. Place the EC2 instances behind an Application Load Balancer (ALB). Configure TCP listeners. Create an AWS Global Accelerator accelerator in front of the ALB. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator.
- D. Place the EC2 instances behind an Amazon CloudFront distribution. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

Answer: B

Explanation:

Question No: 49

A company hosts an application on Amazon EC2 instances behind an Application Load Balancer (ALB).

The company recently experienced a network security breach. A network engineer must collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application.

What is the MOST operationally efficient solution that meets these requirements?

- A. Configure the ALB to store logs in an Amazon S3 bucket. Download the files from Amazon S3, and use a spreadsheet application to analyze the logs.
- B. Configure the ALB to push logs to Amazon Kinesis Data Streams. Use Amazon Kinesis Data Analytics to analyze the logs.
- C. Configure Amazon Kinesis Data Streams to stream data from the ALB to Amazon OpenSearch Service (Amazon Elasticsearch Service). Use search operations in Amazon OpenSearch Service (Amazon Elasticsearch Service) to analyze the data.
- D. Configure the ALB to store logs in an Amazon S3 bucket. Use Amazon Athena to analyze the logs in Amazon S3.

Answer: D

Explanation:

The most operationally efficient solution to collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application would be to configure the ALB to store logs in an Amazon S3 bucket and use Amazon Athena to analyze the logs in Amazon S3 (Option D). This solution allows for quick and easy analysis of log data without requiring manual download or manipulation of log files.

Question No: 50

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=false`. Associate the new DHCP options set with the VPC.
- D. Create a new DHCP options set with parameter `dns_firewall_fail_open=true`. Associate the new DHCP options set with the VPC.

Answer: B

Explanation:

Question No: 51

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

Question No: 52

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 Logs. Review the ALB logs in CloudWatch Insights to determine which error messages the ALB is receiving.
- B. Configure the Amazon S3 bucket destination. Use Amazon Athena to determine which error messages the ALB is receiving.
- C. Configure the Amazon S3 bucket destination. 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.
- D. Send the logs to Amazon CloudWatch Logs. Use the Amazon Athena CloudWatch Connector to determine which error messages the ALB is receiving.

Answer: A

Explanation:

Question No: 53

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 listeners. Use Bring Your Own IP (BYOIP) from the on-premises network with the NLB.

- B. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listeners. Create an AWS Global Accelerator accelerator in front of the NLB. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator.
- C. Place the EC2 instances behind an Application Load Balancer (ALB). Configure TCP listeners. Create an AWS Global Accelerator accelerator in front of the ALB. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator.
- D. Place the EC2 instances behind an Amazon CloudFront distribution. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

Answer: B

Explanation:

Question No: 54

A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin in an Amazon CloudFront distribution. The company wants to implement a custom authentication system that will provide a token for its authenticated customers.

The web application must ensure that the GET/POST requests come from authenticated customers before it delivers the content. A network engineer must design a solution that gives the web application the ability to identify authorized customers.

What is the MOST operationally efficient solution that meets these requirements?

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

Answer: C

Explanation:

Question No: 55

A company's network engineer is designing a hybrid DNS solution for an AWS Cloud workload.

Individual teams want to manage their own DNS hostnames for their applications in their development environment. The solution must integrate the application-specific hostnames with the centrally managed DNS hostnames from the on-premises network and must provide bidirectional name resolution. The solution also must minimize management overhead.

Which combination of steps should the network engineer take to meet these requirements? (Choose three.)

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

Answer: ABE

Explanation:

Question No: 56

A company has its production VPC (VPC-A) in the eu-west-1 Region in Account 1. VPC-A is attached to a transit gateway (TGW-A) that is connected to an on-premises data center in Dublin, Ireland, by an AWS Direct Connect transit VIF that is configured for an AWS Direct Connect gateway. The company also has a staging VPC (VPC-B) that is attached to another transit gateway (TGW-B) in the eu-west-2 Region in Account 2.

A network engineer must implement connectivity between VPC-B and the on-premises data center in Dublin.

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

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

Answer: BC

Explanation:

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

www.VCEplus.io

Question No: 57

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

The solution must minimize cost and administrative overhead.

Which solution will meet these requirements?

- A. Launch an Amazon EC2 instance in the VPC. Use Traffic Mirroring by specifying the NAT gateway as the source and the EC2 instance as the destination. Analyze the captured traffic by using open-source tools to identify the AWS resources that are generating the suspicious traffic.
- B. Use VPC flow logs. Launch a security information and event management (SIEM) solution in the VPC. Configure the SIEM solution to ingest the VPC flow logs. Run queries on the SIEM solution to identify the AWS resources that are generating the suspicious traffic.
- C. Use VPC flow logs. Publish the flow logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query the flow logs to identify the AWS resources that are generating the suspicious traffic.
- D. Configure the VPC to stream the network traffic directly to an Amazon Kinesis data stream. 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

Explanation:

Question No: 58

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

How should the network engineer set up the Direct Connect connection to meet these requirements?

A. Create one hosted connection. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use one Direct Connect gateway for both VIFs to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.

B. Create one hosted connection. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use two Direct Connect gateways, one for each VIF, to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.

C. Create one dedicated connection. Use a transit VIF to connect to the transit gateway in us-east-1.

Use a private VIF to connect to the VPC in eu-west-1. Use one Direct Connect gateway for both VIFs to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.

D. Create one dedicated connection. Use a transit VIF to connect to the transit gateway in us-east-1.

Use a private VIF to connect to the VPC in eu-west-1. Use two Direct Connect gateways, one for each VIF, to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.

Answer: B

Explanation:

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

Question No: 59

A company uses AWS Direct Connect to connect its corporate network to multiple VPCs in the same AWS account and the same AWS Region. Each VPC uses its own private VIF and its own virtual LAN on the Direct Connect connection. The company has grown and will soon surpass the limit of VPCs and private VIFs for each connection.

What is the MOST scalable way to add VPCs with on-premises connectivity?

A. Provision a new Direct Connect connection to handle the additional VPCs. Use the new connection to connect additional VPCs.

B. Create virtual private gateways for each VPC that is over the service quota. Use AWS Site-to-Site VPN to connect the virtual private gateways to the corporate network.

C. Create a Direct Connect gateway, and add virtual private gateway associations to the VPCs.

Configure a private VIF to connect to the corporate network.

D. Create a transit gateway, and attach the VPCs. Create a Direct Connect gateway, and associate it with the transit gateway. Create a transit VIF to the Direct Connect gateway.

Answer: D

Explanation:

When a company requires connectivity to multiple VPCs over AWS Direct Connect, a scalable solution is to use a transit gateway. A transit gateway is a hub that can interconnect multiple VPCs and VPN connections. The VPCs can communicate with each other over the transit gateway, and on-premises networks can communicate with the VPCs through the Direct Connect gateway. This solution provides a central point of management and simplifies the configuration of network routing.

By associating the Direct Connect gateway with the transit gateway, traffic between the VPCs and the on-premises network can be routed through the Direct Connect connection.

Question No: 60

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.

Question No: 61

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

The company tests the application with a single EC2 instance and does not observe any problems.

However, after production deployment, users report that they can log in but that they cannot use the application. Every new web request restarts the login process.

What should a network engineer do to resolve this issue?

- A. Modify the ALB listener configuration. Edit the rule that forwards traffic to the target group.

Change the rule to enable group-level stickiness. Set the duration to the maximum application session length.

- B. Replace the ALB with a Network Load Balancer. Create a TLS listener. Create a new target group with the protocol type set to TLS Register the EC2 instances. Modify the target group configuration by enabling the stickiness attribute.
- C. Modify the ALB target group configuration by enabling the stickiness attribute. Use an applicationbased cookie. Set the duration to the maximum application session length.
- D. Remove the ALB. Create an Amazon Route 53 rule with a failover routing policy for the application name. Configure ACM to issue certificates for each EC2 instance.

Answer: C

Explanation:

Question No: 62

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 Balancer. Run an intrusion detection system (IDS) appliance from AWS Marketplace on Amazon EC2 for traffic inspection.

D. 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

Question No: 63

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. Provide the Connectivity account ID. 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. Associate a route table with the attachment.

B. 1. In the Production account: Create a resource share in AWS Resource Access Manager for the VPC subnets. Provide the Connectivity account ID. 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. Associate a route table with the attachment.

C. 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the VPC subnets. Provide the Production account ID. 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. Associate a route table with the attachment.

D. 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the transit gateway. 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. 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.

Question No: 64

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 subnets. Create an S3 interface endpoint in the VPC. Modify the application configuration to use the S3 endpoint-specific DNS hostname.
- B. Deploy the EC2 instances in the private subnets. Create a NAT gateway in the VPC. Create default routes in the private subnets to the NAT gateway. Connect to Amazon S3 by using the NAT gateway.
- C. Deploy the EC2 instances in the private subnets. Create an S3 gateway endpoint in the VPC. Specify the route table of the private subnets during endpoint creation to create routes to Amazon S3.
- D. Deploy the EC2 instances in the private subnets. Create an S3 interface endpoint in the VPC.

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.

Question No: 65

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. Define the primary network interface to be created in one of the private subnets. For the second network interface, select one of the public subnets. Choose the BYOIP pool ID as the source of public IP addresses.
- B. Configure the primary network interface in a private subnet in the launch template. 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.
- C. Create an AWS Lambda function to run as a lifecycle hook of the Auto Scaling group when an instance is launching. In the Lambda function, assign a network interface to an AWS Global Accelerator endpoint.
- D. 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.

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.

Question No: 66

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

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

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

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

Answer: A

Explanation:

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

Question No: 67

A media company is implementing a news website for a global audience. The website uses Amazon CloudFront as its content delivery network. The backend runs on Amazon EC2 Windows instances behind an Application Load Balancer (ALB). The instances are part of an Auto Scaling group. The company's customers access the website by using service.example.com as the CloudFront custom domain name. The CloudFront origin points to an ALB that uses service-alb.example.com as the domain name.

The company's security policy requires the traffic to be encrypted in transit at all times between the users and the backend.

Which combination of changes must the company make to meet this security requirement? (Choose three.)

- A. Create a self-signed certificate for service.example.com. Import the certificate into AWS Certificate Manager (ACM). Configure CloudFront to use this imported SSL/TLS certificate. Change the default behavior to redirect HTTP to HTTPS.
- B. Create a certificate for service.example.com by using AWS Certificate Manager (ACM). Configure CloudFront to use this custom SSL/TLS certificate. Change the default behavior to redirect HTTP to HTTPS.
- C. Create a certificate with any domain name by using AWS Certificate Manager (ACM) for the EC2 instances. Configure the backend to use this certificate for its HTTPS listener. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its targets.
Attach the existing Auto Scaling group to this new target group.
- D. Create a public certificate from a third-party certificate provider with any domain name for the EC2 instances. Configure the backend to use this certificate for its HTTPS listener. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its targets.
Attach the existing Auto Scaling group to this new target group.
- E. Create a certificate for service-alb.example.com by using AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the service-alb.example.com ACM certificate. Modify the CloudFront origin to use the HTTPS protocol only. Delete the HTTP listener on the ALB.
- F. Create a self-signed certificate for service-alb.example.com. Import the certificate into AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the imported service-alb.example.com ACM certificate. Modify the CloudFront origin to use the HTTPS protocol only. Delete the HTTP listener on the ALB.

Answer: BDE

Explanation:

Question No: 68

Your organization has a newly installed 1-Gbps AWS Direct Connect connection. You order the crossconnect 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

Explanation:

Question No: 69

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/>

Question No: 70

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

Question No: 71

Your company runs an application for the US market in the us-east-1 AWS region. This application uses proprietary TCP and UDP protocols on Amazon Elastic Compute Cloud (EC2) instances. End users run a real-time, front-end application on their local PCs. This front-end application knows the DNS hostname of the service.

You must prepare the system for global expansion. The end users must access the application with lowest latency.

How should you use AWS services to meet these requirements?

- A. Register the IP addresses of the service hosts as "A" records with latency-based routing policy in Amazon Route 53, and set a Route 53 health check for these hosts.
- B. Set the Elastic Load Balancing (ELB) load balancer in front of the hosts of the service, and register the ELB name of the main service host as an ALIAS record with a latency-based routing policy in Route 53.
- C. Set Amazon CloudFront in front of the host of the service, and register the CloudFront name of the main service as an ALIAS record in Route 53.
- D. Set the Amazon API gateway in front of the service, and register the API gateway name of the main service as an ALIAS record in Route 53.

Answer: B

Explanation:

Question No: 72

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

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

Protocol: TCP

Port: 80 inbound, nothing outbound

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

Protocol: TCP

Port: 80 inbound, nothing outbound

When you try to browse to the web server, you receive no response.

Which additional step should you take to receive a successful response?

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

Answer: D

Explanation:

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

Question No: 73

An organization launched an IPv6-only web portal to support IPv6-native mobile clients. Front-end instances launch in an Amazon VPC associated with an appropriate IPv6 CIDR. The VPC IPv4 CIDR is fully utilized. A single subnet exists in each of two Availability Zones with appropriately configured IPv6 CIDR associations. Auto Scaling is properly configured, and no Elastic Load Balancing is used.

Customers say the service is unavailable during peak load times. The network engineer attempts to launch an instance manually and receives the following message: "There are not enough free addresses in subnet 'subnet-12345677' to satisfy the requested number of instances." What action will resolve the availability problem?

- A. Create a new subnet using a VPC secondary IPv6 CIDR, and associate an IPv6 CIDR. Include the new subnet in the Auto Scaling group.
- B. Create a new subnet using a VPC secondary IPv4 CIDR, and associate an IPv6 CIDR. Include the new subnet in the Auto Scaling group.
- C. Resize the IPv6 CIDR on each of the existing subnets. Modify the Auto Scaling group maximum number of instances.
- D. Add a secondary IPv4 CIDR to the Amazon VPC. Assign secondary IPv4 address space to each of the existing subnets.

Answer: B

Explanation:

Question No: 74

An organization is replacing a tape backup system with a storage gateway. there is currently no connectivity to AWS. Initial testing is needed.

What connection option should the organization use to get up and running at minimal cost?

- A. Use an internet connection.
- B. Set up an AWS VPN connection.
- C. Provision an AWS Direct Connection private virtual interface.
- D. Provision a Direct Connect public virtual interface.

Answer: A

Explanation:

Question No: 75

All IP addresses within a 10.0.0.0/16 VPC are fully utilized with application servers across two Availability Zones. The application servers need to send frequent UDP probes to a single central authentication server on the Internet to confirm that is running up-to-date packages. The network is designed for application servers to use a single NAT gateway for internal access. Testing reveals that a few of the servers are unable to communicate with the authentication server.

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

Answer: C

Explanation:

Ref: <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>"A NAT gateway can support up to 55,000 simultaneous connections to each unique destination. This limit also applies if you create approximately 900 connections per second to a single destination (about 55,000 connections per minute). If the destination IP address, the destination port, or the protocol (TCP/UDP/ICMP) changes, you can create an additional 55,000 connections. For more than 55,000 connections, there is an increased chance of connection errors due to port allocation errors.

These errors can be monitored by viewing the ErrorPortAllocation CloudWatch metric for your NAT gateway. For more information, see Monitoring NAT Gateways Using Amazon CloudWatch."

Question No: 76

An organization is using a VPC endpoint for Amazon S3. When the security group rules for a set of instances were initially configured, access was restricted to allow traffic only to the IP addresses of the Amazon S3 API endpoints in the region from the published JSON file. The application was working properly, but now is logging a growing number of timeouts when connecting with Amazon S3. No internet gateway is configured for the VPC.

Which solution will fix the connectivity failures with the LEAST amount of effort?

- A. Create a Lambda function to update the security group based on AmazonIPSpaceChanged notifications.
- B. Update the VPC routing to direct Amazon S3 prefix-list traffic to the VPC endpoint using the route table APIs.
- C. Update the application server's outbound security group to use the prefix-list for Amazon S3 in the same region.
- D. Create an additional VPC endpoint for Amazon S3 in the same route table to scale the concurrent connections to Amazon.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/subscribe-to-aws-public-ip-address-changes-via-amazon-sns/>

Question No: 77

A bank built a new version of its banking application in AWS using containers that content to an onpremises 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 onpremises application version and the rest of the traffic to the new AWS based version.
- B. Use a Classic Load Balancer for the new application. Route all traffic to the new application by using an Elastic Load Balancing (ELB) load balancer DNS. Define a user-agent-based rule on the backend servers to redirect earlier clients to the on-premises application.
- C. Use an Application Load Balancer for the new application. Register both the new and earlier applications as separate target groups and use path-based routing to route traffic based on the application version.
- D. Use an Application Load Balancer for the new application. Register both the new and earlier application backends as separate target groups. Use header-based routing to route traffic based on the application version.

Answer: D

Explanation:

Question No: 78

A company is deploying a non-web application on an AWS load balancer. All targets are servers located on-premises that can be accessed by using AWS Direct Connect. The company wants to ensure that the source IP addresses of clients connecting to the application are passed all the way to the end server.

How can this requirement be achieved?

- A. Use a Network Load Balancer to automatically preserve the source IP address.
- B. Use a Network Load Balancer and enable the X-Forwarded-For attribute.
- C. Use a Network Load Balancer and enable the ProxyProtocol v2 attribute.
- D. Use an Application Load Balancer to automatically preserve the source IP address in the XForwarded- For header.

Answer: C

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-targetgroups.html#proxy-protocol>

Question No: 79

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

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

AWSTemplateFormation Version: 2010-09-09

Parameters:

Originating VPCId:

Type: String

RemoteVPCId:

Type: String

RemoteVPCAccountId:

Type: String

Resources: newVPCPeeringConnection:

Type: 'AWS::EC2::VPCPeeringConnection'

Properties:

VpcId: !Ref OriginatingVPCId

PeerVpcId: !Ref RemoteVPCId

PeerOwnerId: !Ref RemoteVPCAccountId

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

- A. Resources:NewEC2SecurityGroup:Type: AWS::EC2::SecurityGroup
- B. Resources:NetworkInterfaceToRemoteVPC:Type: "AWS::EC2NetworkInterface"
- C. Resources:newEC2Route:Type: AWS::EC2::Route
- D. Resources:VPCGatewayToRemoteVPC:Type: "AWS::EC2::VPCGatewayAttachment"
- E. Resources:newVPCPeeringConnection:Type: 'AWS::EC2VPCPeeringConnection'PeerRoleArn: !Ref PeerRoleArn

Answer: C,E

Explanation:

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

Question No: 80

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

Answer: C

Explanation:

www.VCEplus.io