ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 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) duster 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-target- groups.html#target-group-protocol-versionhttps://docs.aws.amazon.com/prescriptive- guidance/latest/patterns/deploy-a-grpc-based-application-on-an-amazon-eks-cluster-and-access-it- with-an-application-load-balancer.html

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

Question: 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 VPCreate 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: 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: 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/whats- new/2021/09/application-load-balancer-aws-privatelink-static-ip-addresses-network-load-balancer/

Question: 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: 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: 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 VPAdd 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: 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-network- firewall-logs-using-amazon-opensearch-service-part-1/

Question: 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 VPUpdate 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: 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 Auto Scaling 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: 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: 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: 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: 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: 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-and- amazon-route-53/

Question: 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-aws- direct-connect-connections/

Question: 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: 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 us- east-1, the private WAN connection must be used to minimize costs on AWS. A network engineer has configured each transit gateway association on the Direct Connect gateway to advertise VPC-specific CIDR IP prefixes only from the local Region. The routes toward the other Region must be learned through BGP from the routers in the other data center in the original, non-aggregated form.

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

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

A. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connection. 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: 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: 22

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

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

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

A. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the outbound endpoint.

B. Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.

C. Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint.

D. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.

E. Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver.

F. Configure the on-premises DNS resolver to forward aws.example.com queries to the IP addresses of the outbound endpoint.

Answer: BCE

Explanation:

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

Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint. (Option C)

Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint. (Option B)

Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver. (Option E)

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

Question: 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: 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: 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 VPRemove 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: 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: 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 onpremises 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-cloudwatch- events.html

To receive notification each time a new route is advertised to AWS from on premises over Direct Connect, a network engineer should onboard Transit Gateway Network Manager to Amazon CloudWatch Logs Insights and use Amazon EventBridge (Amazon CloudWatch Events) to send notifications when routes change (Option B). This solution allows for real-time monitoring of route changes and automatic notification when new routes are advertised.

Question: 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: 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: 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: 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: 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: 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: 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: 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: 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: 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: 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 pkt- srcaddr 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: 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: 40

A company's network engineer is designing an active-passive connection to AWS from two onpremises data centers. The company has set up AWS Direct Connect connections between the onpremises 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: 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: 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: 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: 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: 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 VPCreate 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: 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: 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: 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 NLUse Bring Your Own IP (BYOIP) from the onpremises 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: 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: 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: 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:

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

An IPsec VPN connection over the transit VIF can encrypt traffic between the on-premises network and AWS without using public IP

addresses or the internet2. A VPC endpoint for Amazon S3 can enable private access to S3 buckets within the same region. HTTPS can provide additional encryption for communication.

Question: 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: B

Explanation:

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

Question: 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 NLUse Bring Your Own IP (BYOIP) from the onpremises 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: 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: 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: 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:

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

Question: 57

A company has deployed an application in a VPC that uses a NAT gateway for outbound traffic to the internet. A network engineer notices a large quantity of suspicious network traffic that is traveling from the VPC over the internet to IP addresses that are

included on a deny list. The network engineer must implement a solution to determine which AWS resources are generating the suspicious traffic. The solution must minimize cost and administrative overhead.

Which solution will meet these requirements?

A. Launch an Amazon EC2 instance in the 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: 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: 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 onpremises 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: 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: 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 application**based 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: 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: 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: 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 VPSpecify die 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: 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: 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: 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: 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: 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: 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 the 0.02 * OUT, ergo VPC peering will be cheaper

Question: 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: 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/knowledge- center/resolve-connection-sg-acl-inbound/

Question: 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: 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: 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: 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: 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: 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 X- Forwarded-For header.

Answer: C

Explanation:

https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target- groups.html#proxy-protocol

Question: 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 VCId:

Type: String

RemoteVPCId:

Type: String

RemoteVPCAccountId:

Type: String

Resources:

newVPCPeeringConnection:

Type: 'AWS::EC2::VPCPeeringConnection'

Properties:

VpcdId: !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: 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:

Question: 81

A company is migrating an existing application to a new AWS account. The company will deploy the application in a single AWS Region by using one VPC and multiple Availability Zones. The application will run on Amazon EC2 instances. Each Availability Zone will have several EC2 instances. The EC2 instances will be deployed in private subnets.

The company's clients will connect to the application by using a web browser with the HTTPS protocol. Inbound connections must be distributed across the Availability Zones and EC2 instances. All connections from the same client session must be connected to the same EC2 instance. The company must provide end-to-end encryption for all connections between the clients and the application by using the application SSL certificate.

Which solution will meet these requirements?

A. Create a Network Load Balancer. Create a target group. Set the protocol to TCP and the port to 443 for the target group. Turn on session affinity (sticky sessions). Register the EC2 instances as targets. Create a listener. Set the protocol to TCP and the port to 443 for the listener. Deploy SSL certificates to the EC2 instances.

B. Create an Application Load Balancer. Create a target group. Set the protocol to HTTP and the port to 80 for the target group. Turn on session affinity (sticky sessions) with an application-based cookie policy. Register the EC2 instances as targets. Create an HTTPS listener. Set the default action to forward to the target group. Use AWS Certificate Manager (ACM) to create a certificate for the listener.

C. Create a Network Load Balancer. Create a target group. Set the protocol to TLS and the port to 443 for the target group. Turn on session affinity (sticky sessions). Register the EC2 instances as targets. Create a listener. Set the protocol to TLS and the port to 443 for the listener. Use AWS Certificate Manager (ACM) to create a certificate for the application.

D. Create an Application Load Balancer. Create a target group. Set the protocol to HTTPS and the port to 443 for the target group. Turn on session affinity (sticky sessions) with an application-based cookie policy. Register the EC2 instances as targets. Create an HTTP listener. Set the port to 443 for the listener. Set the default action to forward to the target group.

Answer: A

Explanation:

Question: 82

A company is developing an application in which IoT devices will report measurements to the AWS Cloud. The application will have millions of end users. The company observes that the IoT devices cannot support DNS resolution. The company needs to implement an Amazon EC2 Auto Scaling solution so that the IoT devices can connect to an application endpoint without using DNS.

Which solution will meet these requirements MOST cost-effectively?

A. Use an Application Load Balancer (ALB)-type target group for a Network Load Balancer (NLB). Create an EC2 Auto Scaling group. Attach the Auto Scaling group to the ALB. Set up the IoT devices to connect to the IP addresses of the NLB.

B. Use an AWS Global Accelerator accelerator with an Application Load Balancer (ALB) endpoint. Create an EC2 Auto Scaling group. Attach the Auto Scaling group to the ALSet up the IoT devices to connect to the IP addresses of the accelerator.

C. Use a Network Load Balancer (NLB). Create an EC2 Auto Scaling group. Attach the Auto Scaling group to the NLB. Set up the IoT devices to connect to the IP addresses of the NLB.

D. Use an AWS Global Accelerator accelerator with a Network Load Balancer (NLB) endpoint. Create an EC2 Auto Scaling group. Attach the Auto Scaling group to the NLB. Set up the IoT devices to connect to the IP addresses of the accelerator.

Answer: D

Explanation:

AWS Global Accelerator can provide static IP addresses that the IoT devices can connect to without using DNS2. It can also route traffic over the AWS global network and improve performance and availability for the IoT devices2. An NLB can provide end-to-end encryption for HTTPS traffic by using TLS as a target group protocol and terminating SSL connections at the load balancer level1. An NLB can also support session affinity (sticky sessions) with TCP connections1.

Question: 83

A company has deployed a new web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Amazon EC2 Auto Scaling group. Enterprise customers from around the world will use the application. Employees of these enterprise customers will connect to the application over HTTPS from office locations.

The company must configure firewalls to allow outbound traffic to only approved IP addresses. The employees of the enterprise customers must be able to access the application with the least amount of latency.

Which change should a network engineer make in the infrastructure to meet these requirements?

A. Create a new Network Load Balancer (NLB). Add the ALB as a target of the NLB.

B. Create a new Amazon CloudFront distribution. Set the ALB as the distribution's origin.

C. Create a new accelerator in AWS Global Accelerator. Add the ALB as an accelerator endpoint.

D. Create a new Amazon Route 53 hosted zone. Create a new record to route traffic to the ALB.

_____ Answer: B

Explanation:

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

Question: 84

A company has hundreds of VPCs on AWS. All the VPCs access the public endpoints of Amazon S3 and AWS Systems Manager through NAT gateways. All the traffic from the VPCs to Amazon S3 and Systems Manager travels through the NAT gateways. The company's network engineer must centralize access to these services and must eliminate the need to use public endpoints.

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

A. Create a central egress VPC that has private NAT gateways. Connect all the VPCs to the central egress VPC by using AWS Transit Gateway. Use the private NAT gateways to connect to Amazon S3 and Systems Manager by using private IP addresses.

B. Create a central shared services VPC. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access. Ensure that private DNS is turned off. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway. Create an Amazon Route 53 forwarding rule for each interface VPC endpoint. Associate the forwarding rules

with all the VPCs. Forward DNS queries to the interface VPC endpoints in the shared services VPC.

C. Create a central shared services VPIn the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access. Ensure that private DNS is turned off. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway. Create an Amazon Route 53 private hosted zone with a full service endpoint name for Amazon S3 and Systems Manager. Associate the private hosted zones with all the VPCs. Create an alias record in each private hosted zone with the full AWS service endpoint pointing to the interface VPC endpoint in the shared services VPC.

D. Create a central shared services VPC. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway. Ensure that private DNS is turned on for the interface VPC endpoints and that the transit gateway is created with DNS support turned on.

Answer: B

Explanation:

Interface VPC endpoints enable private connectivity between VPCs and supported AWS services without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection2. Interface VPC endpoints are powered by AWS PrivateLink, a technology that enables private access to AWS services2. Amazon S3 and AWS Systems Manager support interface VPC endpoints2. By turning off private DNS, the interface VPC endpoints can be accessed by using their private IP addresses2. By using Amazon Route 53 forwarding rules, DNS queries can be resolved to the interface VPC endpoints in the shared services VPC3.

Question: 85

A company manages resources across VPCs in multiple AWS Regions. The company needs to connect to the resources by using its internal domain name. A network engineer needs to apply the aws.example.com DNS suffix to all resources.

What must the network engineer do to meet this requirement?

A. Create an Amazon Route 53 private hosted zone for aws.example.com in each Region that has resources. Associate the private hosted zone with that Region's VPC. In the appropriate private hosted zone, create DNS records for the resources in each Region.

B. Create one Amazon Route 53 private hosted zone for aws.example.com. Configure the private hosted zone to allow zone transfers with every VPC.

C. Create one Amazon Route 53 private hosted zone for example.com. Create a single resource record for aws.example.com in the private hosted zone. Apply a multivalue answer routing policy to the record. Add all VPC resources as separate values in the routing policy.

D. Create one Amazon Route 53 private hosted zone for aws.example.com. Associate the private hosted zone with every VPC that has resources. In the private hosted zone, create DNS records for all resources.

Answer: D

Explanation:

Creating one private hosted zone for aws.example.com and associating it with every VPC that has resources would enable DNS resolution for all resources by using their internal domain name. Creating an alias record in each private hosted zone with the full AWS service endpoint pointing to the interface VPC endpoint in the shared services VPC would enable private connectivity to Amazon S3 and AWS Systems Manager without using public endpoints.

Question: 86

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

Which solution meets these requirements?

A. Configure a private hosted zone for each application VPC, and create the requisite records. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver. Associate the application VPC private hosted zones with the egress VPC, and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints.

B. Configure a public hosted zone for each application VPC, and create the requisite records. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver.
Associate the application VPC private hosted zones with the egress VPC. and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints.

C. Configure a private hosted zone for each application VPC, and create the requisite records. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPDefine Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver.

Associate the application VPC private hosted zones with the egress VPand s

Answer: A

Explanation:

Creating a private hosted zone for each application VPC and creating the requisite records would enable end-to-end domain name resolution for the resources. Creating a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC would enable bi-directional DNS resolution between AWS and the existing on-premises environments. Defining Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver would enable DNS queries from AWS resources to on-premises resources. Associating the application VPC private hosted zones with the egress VPC and sharing the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager would enable DNS queries among different VPCs and accounts. Configuring the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints would enable DNS queries from on-premises resources to AWS resources1.

Question: 87

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

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

What should a network engineer do to reduce this latency?

A. Create a new Site-to-Site VPN connection. Set the transit gateway as the target gateway. Enable acceleration on the new Site-to-Site VPN connection. Update the VPN device in the London office with the new connection details.

B. Modify the existing Site-to-Site VPN connection by setting the transit gateway as the target gateway. Enable acceleration on the existing Site-to-Site VPN connection.

C. Create a new transit gateway in the eu-west-2 (London) Region. Peer the new transit gateway with the existing transit gateway. Modify the existing Site-to-Site VPN connection by setting the new transit gateway as the target gateway.

D. Create a new AWS Global Accelerator standard accelerator that has an endpoint of the Site-to-Site VPN connection. Update the

VPN device in the London office with the new connection details.

Answer: A

Explanation:

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

Question: 88

A company has a hybrid cloud environment. The company's data center is connected to the AWS Cloud by an AWS Direct Connect connection. The AWS environment includes VPCs that are connected together in a hub-and-spoke model by a transit gateway. The AWS environment has a transit VIF with a Direct Connect gateway for on-premises connectivity.

The company has a hybrid DNS model. The company has configured Amazon Route 53 Resolver endpoints in the hub VPC to allow bidirectional DNS traffic flow. The company is running a backend application in one of the VPCs.

The company uses a message-oriented architecture and employs Amazon Simple Queue Service

(Amazon SQS) to receive messages from other applications over a private network. A network engineer wants to use an interface VPC endpoint for Amazon SQS for this architecture. Client services must be able to access the endpoint service from on premises and from multiple VPCs within the company's AWS infrastructure.

Which combination of steps should the network engineer take to ensure that the client applications can resolve DNS for the interface endpoint? (Choose three.)

A. Create the interface endpoint for Amazon SQS with the option for private DNS names turned on.

B. Create the interface endpoint for Amazon SQS with the option for private DNS names turned off.

C. Manually create a private hosted zone for sqs.us-east-1.amazonaws.com. Add necessary records that point to the interface endpoint. Associate the private hosted zones with other VPCs.

D. Use the automatically created private hosted zone for sqs.us-east-1.amazonaws.com with previously created necessary records that point to the interface endpoint. Associate the private hosted zones with other VPCs.

E. Access the SQS endpoint by using the public DNS name sqs.us-east-1 amazonaws.com in VPCs and on premises.

F. Access the SQS endpoint by using the private DNS name of the interface endpoint .sqs.us-east- 1.vpce.amazonaws.com in VPCs and on premises.

Answer: ADF

Explanation:

Question: 89

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

Which solution will meet these requirements?

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a custom pattern to monitor the account for changes. Configure the rule to invoke an AWS Lambda function to identify noncompliant resources. Update an Amazon DynamoDB table with the changes that are identified.

B. Create custom metrics from Amazon CloudWatch logs. Use the metrics to invoke an AWS Lambda function to identify noncompliant resources. Update an Amazon DynamoDB table with the changes that are identified.

C. Record the current state of network resources by using AWS Config. Create rules that reflect the desired configuration settings. Set remediation for noncompliant resources.

D. Record the current state of network resources by using AWS Systems Manager Inventory. Use Systems Manager State Manager to enforce the desired configuration settings and to carry out remediation for noncompliant resources.

_____ **Answer: C**

Explanation:

Recording the current state of network resources by using AWS Config would enable auditing and assessment of resource configurations and compliance^. Creating rules that reflect the desired configuration settings would enable evaluation of whether

the network resources comply with network security policies3. Setting remediation for noncompliant resources would enable automatic correction of undesired configurations3.

Question: 90

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

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

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

A. Set up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional traffic load from remote employees and the additional application. Create a link aggregation group (LAG).

B. Deploy an AWS Site-to-Site VPN connection to the application VPC. Configure the on-premises routing for the remote employees to connect to the Site-to-Site VPN connection.

C. Deploy Amazon Workspaces into the application VPInstruct the remote employees to connect to Workspaces.

D. Replace the existing 1 Gbps Direct Connect connection with two new 2 Gbps Direct Connect hosted connections. Create an AWS Client VPN endpoint in the application VPC. Instruct the remote employees to connect to the Client VPN endpoint.

_____ Answer: A

Explanation:

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

Question: 91

A company has a global network and is using transit gateways to connect AWS Regions together. The company finds that two Amazon EC2 instances in different Regions are unable to communicate with each other. A network engineer needs to troubleshoot this connectivity issue.

What should the network engineer do to meet this requirement?

A. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables and in the VPC route tables. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.

B. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables. Verify that the VPC route tables are correct. Use AWS Firewall Manager to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.

C. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables. Verify that the VPC route tables are correct. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.

D. Use VPC Reachability Analyzer to analyze routes in the transit gateway route tables. Verify that the VPC route tables are correct. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.

Answer: C

Explanation:

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

Question: 92

A company has been using an outdated application layer protocol for communication among applications. The company decides not to use this protocol anymore and must migrate all applications to support a new protocol. The old protocol and the new protocol

are TCP-based, but the protocols use different port numbers.

After several months of work, the company has migrated dozens of applications that run on Amazon EC2 instances and in containers. The company believes that all the applications have been migrated, but the company wants to verify this belief. A network engineer needs to verify that no application is still using the old protocol.

Which solution will meet these requirements without causing any downtime?

A. Use Amazon Inspector and its Network Reachability rules package. Wait until the analysis has finished running to find out which EC2 instances are still listening to the old port.

B. Enable Amazon GuardDuty. Use the graphical visualizations to filter for traffic that uses the port of the old protocol. Exclude all internet traffic to filter out occasions when the same port is used as an ephemeral port.

C. Configure VPC flow logs to be delivered into an Amazon S3 bucket. Use Amazon Athena to query the data and to filter for the port number that is used by the old protocol.

D. Inspect all security groups that are assigned to the EC2 instances that host the applications. Remove the port of the old protocol if that port is in the list of allowed ports. Verify that the applications are operating properly after the port is removed from the security groups.

Answer: C

Explanation:

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

Question: 93

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

- Application VPCs must be isolated from each other.

- Bidirectional communication must be allowed between the application VPCs and the on-premises network.
- Bidirectional communication must be allowed between the application VPCs and the shared services VPC.

The network engineer creates the transit gateway with options disabled for default route table association and default route table propagation. The network engineer also creates the VPN attachment for the on-premises network and creates the VPC attachments for the application VPCs and the shared services VPC.

The network engineer must meet all the requirements for the transit gateway by designing a solution that needs the least number of transit gateway route tables.

Which combination of actions should the network engineer perform to accomplish this goal? (Choose two.)

A. Configure a separate transit gateway route table for on premises. Associate the VPN attachment with this transit gateway route table. Propagate all application VPC attachments to this transit gateway route table.

B. Configure a separate transit gateway route table for each application VPC. Associate each application VPC attachment with its respective transit gateway route table. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.

C. Configure a separate transit gateway route table for all application VPCs. Associate all application VPCs with this transit gateway route table. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.

D. Configure a separate transit gateway route table for the shared services VPC. Associate the shared services VPC attachment with this transit gateway route table. Propagate all application VPC attachments to this transit gateway route table.

E. Configure a separate transit gateway route table for on premises and the shared services VPC. Associate the VPN attachment and the shared services VPC attachment with this transit gateway route table. Propagate all application VPC attachments to this transit gateway route table.

Answer: BD

Explanation:

Question: 94

A company has an AWS Site-to-Site VPN connection between its existing VPC and on-premises network. The default DHCP options set is associated with the VPC. The company has an application that is running on an Amazon Linux 2 Amazon EC2 instance in the

VPC. The application must retrieve an Amazon RDS database secret that is stored in AWS Secrets Manager through a private VPC endpoint. An on-premises application provides internal RESTful API service that can be reached by URL (https://api.example.internal). Two on-premises Windows DNS servers provide internal DNS resolution.

The application on the EC2 instance needs to call the internal API service that is deployed in the onpremises environment. When the application on the EC2 instance attempts to call the internal API service by referring to the hostname that is assigned to the service, the call fails. When a network engineer tests the API service call from the same EC2 instance by using the API service's IP address, the call is successful.

What should the network engineer do to resolve this issue and prevent the same problem from affecting other resources in the VPC?

A. Create a new DHCP options set that specifies the on-premises Windows DNS servers. Associate the new DHCP options set with the existing VPC. Reboot the Amazon Linux 2 EC2 instance.

B. Create an Amazon Route 53 Resolver rule. Associate the rule with the VPC. Configure the rule to forward DNS queries to the on-premises Windows DNS servers if the domain name matches example.internal.

C. Modify the local host file in the Amazon Linux 2 EC2 instance in the VPMap the service domain name (api.example.internal) to the IP address of the internal API service.

D. Modify the local /etc/resolv.conf file in the Amazon Linux 2 EC2 instance in the VPC. Change the IP addresses of the name servers in the file to the IP addresses of the company's on-premises Windows DNS servers.

Answer: B

Explanation:

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

Question: 95

A company has several production applications across different accounts in the AWS Cloud. The company operates from the us-east-1 Region only. Only certain partner companies can access the applications. The applications are running on Amazon EC2

instances that are in an Auto Scaling group behind an Application Load Balancer (ALB). The EC2 instances are in private subnets and allow traffic only from the ALB. The ALB is in a public subnet and allows inbound traffic only from partner network IP address ranges over port 80.

When the company adds a new partner, the company must allow the IP address range of the partner network in the security group that is associated with the ALB in each account. A network engineer must implement a solution to centrally manage the partner network IP address ranges.

Which solution will meet these requirements in the MOST operationally efficient manner?

A. Create an Amazon DynamoDB table to maintain all IP address ranges and security groups that need to be updated. Update the DynamoDB table with the new IP address range when the company adds a new partner. Invoke an AWS Lambda function to read new IP address ranges and security groups from the DynamoDB table to update the security groups. Deploy this solution in all accounts.

B. Create a new prefix list. Add all allowed IP address ranges to the prefix list. Use Amazon EventBridge (Amazon CloudWatch Events) rules to invoke an AWS Lambda function to update security groups whenever a new IP address range is added to the prefix list. Deploy this solution in all accounts.

C. Create a new prefix list. Add all allowed IP address ranges to the prefix list. Share the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM). Update security groups to use the prefix list instead of the partner IP address range. Update the prefix list with the new IP address range when the company adds a new partner.

D. Create an Amazon S3 bucket to maintain all IP address ranges and security groups that need to be updated. Update the S3 bucket with the new IP address range when the company adds a new partner. Invoke an AWS Lambda function to read new IP address ranges and security groups from the S3 bucket to update the security groups. Deploy this solution in all accounts.

Answer: C

Explanation:

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

Question: 96

A company is migrating an application from on premises to AWS. The company will host the application on Amazon EC2 instances that are deployed in a single VPC. During the migration period, DNS queries from the EC2 instances must be able to resolve names of on-premises servers. The migration is expected to take 3 months After the 3-month migration period, the resolution of on-premises servers will no longer be needed.

What should a network engineer do to meet these requirements with the LEAST amount of configuration?

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

B. Set up an AWS Direct Connect connection with a private VIF. Deploy an Amazon Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint in the Region that is hosting the VPC.

C. Set up an AWS Client VPN connection between on premises and AWS. Deploy an Amazon Route 53 Resolver inbound endpoint in the VPC.

D. Set up an AWS Direct Connect connection with a public VIF. Deploy an Amazon Route 53 Resolver inbound endpoint in the Region that is hosting the VPC. Use the IP address that is assigned to the endpoint for connectivity to the on-premises DNS servers.

_____ Answer: A

Explanation:

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

Question: 97

A company is hosting an application on Amazon EC2 instances behind an Application Load Balancer.

The instances are in an Amazon EC2 Auto Scaling group. Because of a recent change to a security group, external users cannot access the application.

A network engineer needs to prevent this downtime from happening again. The network engineer must implement a solution that remediates noncompliant changes to security groups.

Which solution will meet these requirements?

A. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuration. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.

B. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.

C. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuration. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.

D. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.

_____ Answer: D

Explanation:

Configuring an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration would enable evaluation of the compliance status of the security groups based on predefined or custom rules3. Creating an AWS Systems Manager Automation runbook to remediate noncompliant security groups would enable automation of the remediation process2. Additionally, configuring AWS Config to trigger the runbook when a noncompliant change is detected would enable timely and consistent remediation of security group changes.

Question: 98

A company is deploying third-party firewall appliances for traffic inspection and NAT capabilities in its VPC. The VPC is configured

with private subnets and public subnets. The company needs to deploy the firewall appliances behind a load balancer.

Which architecture will meet these requirements MOST cost-effectively?

A. Deploy a Gateway Load Balancer with the firewall appliances as targets. Configure the firewall appliances with a single network interface in a private subnet. Use a NAT gateway to send the traffic to the internet after inspection.

B. Deploy a Gateway Load Balancer with the firewall appliances as targets. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subnet. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.

C. Deploy a Network Load Balancer with the firewall appliances as targets. Configure the firewall appliances with a single network interface in a private subnet. Use a NAT gateway to send the traffic to the internet after inspection.

D. Deploy a Network Load Balancer with the firewall appliances as targets. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subnet. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.

Answer: B

Explanation:

Question: 99

A company's AWS architecture consists of several VPCs. The VPCs include a shared services VPC and several application VPCs. The company has established network connectivity from all VPCs to the on-premises DNS servers.

Applications that are deployed in the application VPCs must be able to resolve DNS for internally hosted domains on premises. The applications also must be able to resolve local VPC domain names and domains that are hosted in Amazon Route 53 private hosted zones.

What should a network engineer do to meet these requirements?

A. Create a new Route 53 Resolver inbound endpoint in the shared services VPC. Create forwarding rules for the on-premises hosted domains. Associate the rules with the new Resolver endpoint and each application VPC. Update each application VPC's DHCP

configuration to point DNS resolution to the new Resolver endpoint.

B. Create a new Route 53 Resolver outbound endpoint in the shared services VPC. Create forwarding rules for the on-premises hosted domains. Associate the rules with the new Resolver endpoint and each application VPC.

C. Create a new Route 53 Resolver outbound endpoint in the shared services VPCreate forwarding rules for the on-premises hosted domains. Associate the rules with the new Resolver endpoint and each application VPUpdate each application VPC's DHCP configuration to point DNS resolution to the new Resolver endpoint.

D. Create a new Route 53 Resolver inbound endpoint in the shared services VPC. Create forwarding rules for the on-premises hosted domains. Associate the rules with the new Resolver endpoint and each application VPC.

_____ Answer: B

Explanation:

Creating a new Route 53 Resolver outbound endpoint in the shared services VPC would enable forwarding of DNS queries from the VPC to on-premises1. Creating forwarding rules for the onpremises hosted domains would enable specifying which domain names are forwarded to the onpremises DNS servers2. Associating the rules with the new Resolver endpoint and each application VPC would enable applying the rules to the VPCs2. This solution would not affect the default DNS resolution behavior of Route 53 Resolver for local VPC domain names and domains that are hosted in Route 53 private hosted zones3.

Question: 100

A company uses Amazon Route 53 for its DNS needs. The company's security team wants to update the DNS infrastructure to provide the most recent security posture.

The security team has configured DNS Security Extensions (DNSSEC) for the domain. The security team wants a network engineer to explain who is responsible for the rotation of DNSSEC keys.

Which explanation should the network administrator provide to the security team?

A. AWS rotates the zone-signing key (ZSK). The company rotates the key-signing key (KSK).

B. The company rotates the zone-signing key (ZSK) and the key-signing key (KSK).

C. AWS rotates the AWS Key Management Service (AWS KMS) key and the key-signing key (KSK).

D. The company rotates the AWS Key Management Service (AWS KMS) key. AWS rotates the key-signing key (KSK).

Answer: A

Explanation:

Question: 101

AnyCompany has acquired Example Corp. AnyCompany's infrastructure is all on premises, and Example Corp's infrastructure is completely in the AWS Cloud. The

companies are using AWS Direct Connect with AWS Transit Gateway to establish connectivity between each other.

Example Corp has deployed a new application across two Availability Zones in a VPC with no internet gateway. The CIDR range for the VPC is 10.0.0.0/16. Example

Corp needs to access an application that is deployed on premises by AnyCompany. Because of compliance requirements, Example Corp must access the application

through a limited contiguous block of approved IP addresses (10.1.0.0/24).

A network engineer needs to implement a highly available solution to achieve this goal. The network engineer starts by updating the VPC to add a new CIDR range of

10.1.0.0/24.

What should the network engineer do next to meet the requirements?

A. In each Availability Zone in the VPC, create a subnet that uses part of the allowed IP address range. Create a public NAT Sateway in each of the new

subnets. Update the route tables that are associated with other subnets to route application traffic to the public NAT gateway in the corresponding Availability

Zone. Add a route to the route table that is associated with the subnets of the public NAT gateways to send traffic destined for the application to the transit

gateway.

B. In each Availability Zone in the VPC, create a subnet that uses part of the allowed IP address range. Create a private NAT gateway in each of the new

subnets. Update the route tables that are associated with other subnets to route application traffic to the private NAT gateway in the corresponding

Availability Zone. Add a route to the route table that is associated with the subnets of the private

NAT gateways to send traffic destined for the application to

the transit gateway.

C. In the VPC, create a subnet that uses the allowed IP address range. Create a private NAT gateway in the new subnet. Update the route tables that are associated with other subnets to route application traffic to the private NAT gateway. Add a route to the route table that is associated with the subnet of the private NAT gateway to send traffic destined for the application to the transit gateway.

D. In the VPC, create a subnet that uses the allowed IP address range. Create a public NAT gateway in the new subnet. Update the route tables that are associated with other subnets to route application traffic to the public NAT gateway. Add a route to the route table that is associated with the subnet of the public NAT gateway to send traffic destined for the application to the transit gateway.

_____ Answer: B

Explanation:

The correct answer is B. In each Availability Zone in the VPC, create a subnet that uses part of the allowed IP address range. Create a private NAT gateway in each of the new subnets. Update the route tables that are associated with other subnets to route application traffic to the private NAT gateway in the corresponding Availability Zone. Add a route to the route table that is associated with the subnets of the private NAT gateways to send traffic destined for the application to the transit gateway.

This solution meets the requirements because:

- It uses a private NAT gateway, which can route traffic to other VPCs or on-premises networks through a transit gateway or a virtual private gateway1.
- It creates a subnet in each Availability Zone that uses part of the approved IP address range, which ensures high availability and compliance.
- It updates the route tables to send traffic from the other subnets to the private NAT gateway in the same Availability Zone, which reduces latency and improves performance.
- It adds a route to the route table of the private NAT gateway subnets to send traffic destined for the application to the transit gateway, which enables connectivity to the on-premises network.

The other options are incorrect because:

- Option A uses a public NAT gateway, which is not necessary for connecting to other VPCs or on-premises networks. A

public NAT gateway also requires an elastic IP address, which is not part of the approved IP address range.

- Option C creates only one subnet and one private NAT gateway, which does not provide high availability across multiple Availability Zones.
- Option D uses a public NAT gateway, which is not necessary for connecting to other VPCs or on-premises networks. A public NAT gateway also requires an elastic IP address, which is not part of the approved IP address range. Additionally, option D creates only one subnet and one public NAT gateway, which does not provide high availability across multiple Availability Zones.

Question: 102

A network engineer is working on a large migration effort from an on-premises data center to an AWS Control Tower based multi-account environment. The environment

has a transit gateway that is deployed to a central network services account. The central network services account has been shared with an organization in AWS

Organizations through AWS Resource Access Manager (AWS RAM).

A shared services account also exists in the environment. The shared services account hosts workloads that need to be shared with the entire organization.

The network engineer needs to create a solution to automate the deployment of common network components across the environment. The solution must provision a

VPC for application workloads to each new and existing member account. The VPCs must be connected to the transit gateway in the central network services account.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Select THREE.)

A. Deploy an AWS Lambda function to the shared services account. Program the Lambda function to assume a role in the new and existing member accounts

to provision the necessary network infrastructure.

B. Update the existing accounts with an Account Factory Customization (AFC). Select the same AFC when provisioning new accounts.

C. Create an AWS CloudFormation template that describes the infrastructure that needs to be created in each account. Upload the template as an AWS

Service Catalog product to the shared services account.

D. Deploy an Amazon EventBridge rule on a default event bus in the shared services account. Configure the EventBridge rule to react to AWS Control Tower

CreateManagedAccount lifecycle events and to invoke the AWS Lambda function.

E. Create an AWSControlTowerBlueprintAccess role in the shared services account.

F. Create an AWSControlTowerBlueprintAccess role in each member account.

Answer: ACD

Explanation:

The correct answer is A, C, and D. These steps will meet the requirements with the least operational overhead because:

- Step A will deploy an AWS Lambda function to the shared services account that can automate the network infrastructure provisioning in each member account by assuming a role with the necessary permissions.
- Step C will create an AWS CloudFormation template that describes the VPC and the transit gateway attachment for each account. This template can be uploaded as an AWS Service Catalog product to the shared services account, which can be used by the AWS Lambda function to create the network resources in each member account.
- Step D will deploy an Amazon EventBridge rule on a default event bus in the shared services account that can react to AWS Control Tower lifecycle events, such as creating a new managed account. This rule can invoke the AWS Lambda function to provision the network infrastructure in the new account.

The other steps are incorrect because:

- Step B will update the existing accounts with an Account Factory Customization (AFC), which is a feature of AWS Control Tower that allows you to customize the account creation process with AWS CloudFormation templates. However, this step will not automate the network infrastructure provisioning for the existing accounts, as it only applies to the new accounts created through the

Account Factory. Moreover, this step will require additional operational overhead to maintain the AFC templates and products.

- Step E will create an AWSControlTowerBlueprintAccess role in the shared services account, which is a role that allows AWS Control Tower to access the AWS Service Catalog products in the shared services account. However, this step is not necessary for the automation solution, as the AWS Lambda function can access the AWS Service Catalog products directly without using this role.
- Step F will create an AWSControlTowerBlueprintAccess role in each member account, which is a role that allows AWS Control Tower to access the AWS Service Catalog products in the member accounts. However, this step is not necessary for the automation solution, as the AWS Lambda function can access the AWS Service Catalog products in the shared services account without using this role.

Question: 103

A company has a total of 30 VPCs. Three AWS Regions each contain 10 VPCs. The company has attached the VPCs in each Region to a transit gateway in that Region. The company also

has set up inter-Region peering connections between the transit gateways.

The company wants to use AWS Direct Connect to provide access from its on-premises location for only four VPCs across the three Regions. The company has provisioned four Direct

Connect connections at two Direct Connect locations.

Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

A. Create four virtual private gateways. Attach the virtual private gateways to the four VPCs.

B. Create a Direct Connect gateway. Associate the four virtual private gateways with the Direct Connect gateway.

C. Create four transit VIFs on each Direct Connect connection. Associate the transit VIFs with the Direct Connect gateway.

D. Create four transit VIFs on each Direct Connect connection. Associate the transit VIFs with the four virtual private gateways.

E. Create four private VIFs on each Direct Connect connection to the Direct Connect gateway.

F. Create an association between the Direct Connect gateway and the transit gateways.

Answer: BCF

Explanation:

To connect to multiple VPCs across different Regions using Direct Connect, the best option is to use a Direct Connect gateway and transit gateways. A Direct Connect gateway allows you to associate multiple virtual private gateways and transit gateways with the same Direct Connect connection. A transit gateway acts as a network hub that connects multiple VPCs and on-premises networks. By creating inter-Region peering connections between the transit gateways, you can enable crossRegion communication. Therefore, the steps are:

- Create four virtual private gateways and attach them to the four VPCs that need access from the on-premises location.
- Create a Direct Connect gateway and associate it with the four virtual private gateways.
- Create four transit VIFs on each Direct Connect connection and associate them with the Direct Connect gateway. A transit VIF allows you to connect to a Direct Connect gateway using a private ASN.
- Create an association between the Direct Connect gateway and the transit gateways in each Region. This will enable the on-premises location to access the VPCs that are attached to the transit gateways.

Question: 104

A company's VPC has Amazon EC2 instances that are communicating with AWS services over the public internet. The company needs to change the connectivity so that the communication
does not occur over the public intemet.

The company deploys AWS PrivateLink endpoints in the VPC. After the deployment of the PrivateLink endpoints, the EC2 instances can no longer communicate at all with the required AWS services.

Which combination of steps should a network engineer take to restore communication with the AWS services? (Select TWO.)

A. In the VPC route table, add a route that has the PrivateLink endpoints as the destination.

B. Ensure that the enableDnsSupport attribute is set to True for the VPC. Ensure that each VPC endpoint has DNS support enabled.

C. Ensure that the VPC endpoint policy allows communication.

D. Create an Amazon Route 53 public hosted zone for all services.

E. Create an Amazon Route 53 private hosted zone that includes a custom name for each service.

Answer: BC

Explanation:

To use AWS PrivateLink, you need to create interface type VPC endpoints for the services that you want to access privately from your VPC1. These endpoints appear as elastic network interfaces (ENIs) with private IPs in your subnets2. To enable DNS resolution for these endpoints, you need to set the enableDnsSupport attribute to True for your VPC, and enable DNS support for each endpoint3. You also need to ensure that the VPC endpoint policy allows communication between your VPC and the service4. You do not need to create any route table entries or Route 53 hosted zones for the endpoints, as they are not required for PrivateLink5.

AWS Private Li nk FAQs - Amazon Web Services 2: AWS PrivateLink and service endpoint - Amazon EC2 Overview and Networking Introduction for Telecom Companies 3: VPC Endpoints: Secure and Direct Access to AWS Services 4: AWS PrivateLink and service endpoint - Amazon EC2 Overview and Networking Introduction for Telecom Companies 5: AWS Private Link vs VPC Endpoint - Stack Overflow

Question: 105

A company is planning to migrate an internal application to the AWS Cloud. The application will run on Amazon EC2 instances in one VPC. Users will access the application from the company's on-premises data center through AWS VPN or AWS Direct Connect. Users will use private domain names for the application endpoint from a domain name that is reserved explicitly for use in the AWS Cloud.

Each EC2 instance must have automatic failover to another EC2 instance in the same AWS account and the same VPC. A

network engineer must design a DNS solution that will not expose the application to the internet.

Which solution will meet these requirements?

A. Assign public IP addresses to the EC2 instances. Create an Amazon Route 53 private hosted zone for the AWS reserved domain name. Associate the private hosted zone with the VPC. Create a Route 53 Resolver outbound endpoint. Configure conditional forwarding in the on-premises DNS resolvers to forward all DNS queries for the AWS domain to the outbound endpoint IP address for Route 53 Resolver. In the private hosted zone, configure primary and failover records that point to the public IP addresses of the EC2 instances. Create an Amazon CloudWatch metric and alarm to monitor the application's health. Set up a health check on the alarm for the primary application endpoint.

B. Place the EC2 instances in private subnets. Create an Amazon Route 53 public hosted zone for the AWS reserved domain name. Associate the public hosted zone with the VPC. Create a Route 53 Resolver inbound endpoint. Configure conditional forwarding in the on-premises DNS resolvers to forward all DNS queries for the AWS domain to the inbound endpoint IP address for Route 53 Resolver. In the public hosted zone, configure primary and failover records that point to the IP addresses of the EC2 instances.
Create an Amazon CloudWatch metric and alarm to monitor the application's health. Set up a health check on the alarm for the primary application endpoint.

C. Place the EC2 instances in private subnets. Create an Amazon Route 53 private hosted zone for the AWS reserved domain name. Associate the private hosted zone with the VPC. Create a Route 53 Resolver inbound endpoint. Configure conditional forwarding in the on-premises DNS resolvers to forward all DNS queries for the AWS domain to the inbound endpoint IP address for Route 53 Resolver. In the private hosted zone, configure primary and failover records that point to the IP addresses of the EC2 instances.
Create an Amazon CloudWatch metric and alarm to monitor the application's health. Set up a health check on the alarm for the primary application endpoint.

D. Place the EC2 instances in private subnets. Create an Amazon Route 53 private hosted zone for the AWS reserved domain name. Associate the private hosted zone with the

VPC. Create a Route 53 Resolver inbound endpoint. Configure conditional forwarding in the on-premises DNS resolvers to forward all DNS queries for the AWS domain to the

inbound endpoint IP address for Route 53 Resolver. In the private hosted zone, configure primary and failover records that point to the IP addresses of the EC2 instances. Set

up Route 53 health checks on the private IP addresses of the EC2 instances.

Answer: C

Explanation:

The correct solution is to use a Route 53 private hosted zone and a Route 53 Resolver inbound endpoint. A private hosted zone allows you to use private domain names for your internal AWS resources without exposing them to the internet. A Route 53 Resolver inbound endpoint enables DNS queries from your on-premises network to be forwarded to your VPC. By configuring conditional forwarding on your on-premises DNS resolvers, you can ensure that only the queries for the AWS reserved domain name are sent to the inbound endpoint. In the private hosted zone, you can create primary and failover records that point to the IP addresses of the EC2 instances. These records will automatically switch to the failover instance if the primary instance becomes unhealthy. You can use CloudWatch metrics and alarms to monitor the application's health and trigger the health check for the primary endpoint.

The other options are not correct because they either expose the application to the internet or use a public hosted zone, which is not suitable for internal applications. Option A assigns public IP addresses to the EC2 instances, which makes them accessible from the internet. Option B uses a public hosted zone, which requires the EC2 instances to have public IP addresses or elastic IP addresses. Option D does not set up a health check on the alarm for the primary endpoint, which is required for the failover mechanism to work.

Question: 106

A company is using an Amazon CloudFront distribution that is configured with an Application Load Balancer (ALB) as an origin. A network engineer needs to implement a solution that requires

all inbound traffic to the ALB to come from CloudFront. The network engineer must implement the solution at the network layer rather than in the application.

Which solution will meet these requirements in the MOST operationally efficient way?

A. Add an inbound rule to the ALB's security group to allow the AWS managed prefix list for CloudFront.

B. Add an inbound rule to the network ACLs that are associated with the ALB's subnets. Use the AWS managed prefix list for CloudFront as the source in the rule.

C. Configure CloudFront to add a custom HTTP header to the requests that CloudFront sends to the ALB.

D. Associate an AWS WAF web ACL with the ALB. Configure the AWS WAF rules to allow traffic from the CloudFront IP set. Automatically update the CloudFront IP set by using an AWS Lambda function.

Answer: A

Explanation:

The most operationally efficient way to restrict inbound traffic to the ALB to come from CloudFront is to use the AWS managed prefix list for CloudFront. A prefix list is a collection of CIDR blocks that can be used to configure security groups and network ACLs. AWS provides a managed prefix list for CloudFront that is automatically updated when CloudFront IP ranges change. By adding an inbound rule to the ALB's security group to allow the AWS managed prefix list for CloudFront, the network engineer can ensure that only CloudFront can access the ALB at the network layer. This solution does not require any additional configuration or maintenance. Option B is less efficient because network ACLs are stateless and require rules for both inbound and outbound traffic. Option C is not a network layer solution, but an application layer solution that requires the ALB to inspect the HTTP headers and reject requests that do not have the custom header. Option D is also not a network layer solution, but a web layer solution that requires AWS WAF to filter the traffic based on the CloudFront IP set. This solution also requires an AWS Lambda function to update the CloudFront IP set, which adds complexity and cost

Question: 107

A company has two business units (BUs). The company operates in the us-east-1 Region and the us- west-1 Region. The company plans to extend to more Regions in the future. Each BU has

a VPC in each Region. Each Region has a transit gateway with the BU VPCs attached. The transit gateways in both Regions are peered.

The company will create several more BUs in the future and will need to isolate some of the BUs from the other BUs. The company wants to migrate to an architecture to incorporate more
Regions and BUs.

Which solution will meet these requirements with the MOST operational efficiency?

A. Create a new transit gateway for each new BU in each Region. Peer the new transit gateways with the existing transit gateways. Update the route tables to control traffic between BUs.

B. Create an AWS Cloud WAN core network with an edge location in both Regions. Configure a segment for each BU with VPC attachments to the new BU VPCs. Use segment actions to control traffic between segments.

C. Create an AWS Cloud WAN core network with an edge location in both Regions. Configure a segment for each BU with VPC attachments to the new BU VPCs. Configure the segments to isolate attachments to control traffic between segments.

D. Attach new VPCs to the existing transit gateways. Update route tables to control traffic between BUs.

Answer: C

Explanation:

The correct solution is to use AWS Cloud WAN, which is a new service that simplifies the management of global networks. AWS Cloud WAN allows you to create a core network that connects your AWS Regions and on-premises locations. You can then create segments for each BU and attach their VPCs to the segments. By configuring the segments to isolate attachments, you can prevent traffic from flowing between different BUs. This way, you can achieve network isolation and scalability without creating multiple transit gateways and peering connections. You can also use segment actions to apply routing and security policies to the traffic within and across segments.

Question: 108

A company has established connectivity between its on-premises data center in Paris, France, and the AWS Cloud by using an AWS

Direct Connect connection. The company uses a transit VIF that connects the Direct Connect connection with a transit gateway that is hosted in the Europe (Paris) Region. The company hosts workloads in private subnets in several VPCs that are attached to the transit gateway.

The company recently acquired another corporation that hosts workloads on premises in an office building in Tokyo, Japan. The company needs to migrate the workloads from the Tokyo office to AWS. These workloads must have access to the company's existing workloads in Paris. The company also must establish connectivity between the Tokyo office building and the Paris data center.

In the Asia Pacific (Tokyo) Region, the company creates a new VPC with private subnets for migration of the workloads. The workload migration must be completed in 5 days. The workloads cannot be directly accessible from the internet.

Which set of steps should a network engineer take to meet these requirements?

A.

1. Create public subnets in the Tokyo VPC to migrate the workloads into.
2. Configure an internet gateway for the Tokyo office to reach the Tokyo VPC.
3. Configure security groups on the Tokyo workloads to only allow traffic from the Tokyo office and the Paris workloads.
4. Create peering connections between the Tokyo VPC and the Paris VPCs.
5. Configure a VPN connection between the Paris data center and the Tokyo office by using existing routers.

B.

1. Configure a transit gateway in the Asia Pacific (Tokyo) Region. Associate this transit gateway with the Tokyo VPC.
2. Create peering connections between the Tokyo transit gateway and the Paris transit gateway.
3. Set up a new Direct Connect connection from the Tokyo office to the Tokyo transit gateway.
4. Configure routing on both transit gateways to allow data to flow between sites and the VPCs.

C.

1. Configure a transit gateway in the Asia Pacific (Tokyo) Region. Associate this transit gateway with the Tokyo VPC.
2. Create peering connections between the Tokyo transit gateway and the Paris transit gateway.
3. Configure an AWS Site-to-Site VPN connection from the Tokyo office. Set the Tokyo transit gateway as the target.

4. Configure routing on both transit gateways to allow data to flow between sites and the VPCs.

D.

1. Configure an AWS Site-to-Site VPN connection from the Tokyo office to the Paris transit gateway.
2. Create an association between the Paris transit gateway and the Tokyo VPC.
3. Configure routing on the Paris transit gateway to allow data to flow between sites and the VPCs.

Answer: C

Explanation:

Option C is the best solution because it allows the company to use transit gateways to connect the VPCs in different regions and the on-premises sites. Transit gateways support inter-region peering and VPN attachments, which enable secure and scalable connectivity. Option A is not valid because public subnets are not suitable for workloads that cannot be directly accessible from the internet. Option B is not valid because Direct Connect connections take longer than 5 days to provision.

Question: 109

A company needs to manage Amazon EC2 instances through command line interfaces for Linux hosts and Windows hosts. The EC2 instances are deployed in an environment in which there is no route to the internet. The company must implement role-based access control for management of the instances. The company has a standalone on-premises environment.

Which approach will meet these requirements with the LEAST maintenance overhead?

A. Set up an AWS Direct Connect connection between the on-premises environment and the VPC where the instances are deployed. Configure routing, security groups, and ACLs.

Connect to the instances by using the Direct Connect connection.

B. Deploy and configure AWS Systems Manager Agent (SSM Agent) on each instance. Deploy VPC endpoints for Systems Manager Session Manager. Connect to the instances by

using Session Manager.

C. Establish an AWS Site-to-Site VPN connection between the on-premises environment and the VPC where the instances are deployed. Configure routing, security groups, and ACLs. Connect to the instances by using the Site-to-Site VPN connection.

D. Deploy an appliance to the VPC where the instances are deployed. Assign a public IP address to the appliance. Configure security groups and ACLs. Connect to the instances by using the appliance as an intermediary.

Answer: B

Explanation:
The correct approach is to use AWS Systems Manager Session Manager, which allows you to manage your EC2 instances through a secure and browser-based interface. By deploying and configuring SSM Agent on each instance, you can enable Session Manager to communicate with the instances. By deploying VPC endpoints for Session Manager, you can enable the instances to connect to the AWS service without requiring an internet gateway, NAT device, or VPN connection. You can also use IAM policies and SSM documents to implement role-based access control for managing the instances. This approach has the least maintenance overhead, as it does not require any additional infrastructure or configuration.

Question: 110

A company has workloads that run in a VPC. The workloads access Amazon S3 by using an S3 gateway endpoint. The company also has on-premises workloads that need to access Amazon S3 privately over a VPN connection. The company has established the VPN connection to the VPC.

Which solution will provide connectivity to Amazon S3 from the VPC workloads and the on-premises workloads in the MOST operationally efficient way?

A. Deploy a proxy fleet of Amazon EC2 instances in the VPC behind an Application Load Balancer (ALB). Configure the on-premises workloads to use the ALB as the proxy server to connect to Amazon S3. Configure the proxy fleet to use the S3 gateway endpoint to connect to Amazon S3.

B. Delete the S3 gateway endpoint. Create an S3 interface endpoint. Deploy a proxy fleet of Amazon EC2 instances in the VPC behind an Application Load Balancer (ALB).

Configure the on-premises workloads to use the ALB as the proxy server to connect to Amazon S3. Configure the proxy fleet and the VPC workloads to use the S3 interface

endpoint to connect to Amazon S3.

C. Create an S3 interface endpoint. Configure an on-premises DNS resolver to resolve the S3 DNS names to the private IP addresses of the S3 interface endpoint. Use the S3

interface endpoint to access Amazon S3. Continue to use the S3 gateway endpoint for the VPC workloads to access Amazon S3.

D. Set up an AWS Direct Connect connection. Create a public VIF. Configure on-premises routing to route the S3 traffic over the public VIF. Make no changes to the on-premises

workloads. Continue to use the S3 gateway endpoint for the VPC workloads to access Amazon S3.

Answer: C

Explanation:

The correct solution is to use an S3 interface endpoint and an on-premises DNS resolver. An S3 interface endpoint allows you to access Amazon S3 using private IP addresses within your VPC. An on-premises DNS resolver can be configured to forward the DNS queries for the S3 domain names to the S3 interface endpoint, so that the on-premises workloads can access Amazon S3 privately over the VPN connection. This solution is operationally efficient, as it does not require any additional infrastructure or changes to the existing workloads. The VPC workloads can continue to use the S3 gateway endpoint, which provides lower latency and higher throughput than the S3 interface endpoint.

Question: 111

A company has many application VPCs that use AWS Site-to-Site VPN connections for connectivity to an on-premises location. The company's network team wants to gradually migrate to AWS Transit Gateway to provide VPC-to-VPC connectivity.

The network team sets up a transit gateway that uses equal-cost multi-path (ECMP) routing. The network team attaches two temporary VPCs to the transit gateway for testing. The test VPCs contain Amazon EC2 instances to confirm connectivity over the transit gateway between the on-premises location and the VPCs. The network team creates two new Site-to-Site VPN connections to the transit gateway.

During testing, the network team cannot reach the required bandwidth of 2.5 Gbps over the pair of new Site-to-Site VPN

connections.

Which combination of steps should the network team take to improve bandwidth performance and minimize network congestion? (Select THREE.)

A. Enable acceleration for the existing Site-to-Site VPN connections to the transit gateway.

B. Create new accelerated Site-to-Site VPN connections to the transit gateway.

C. Advertise the on-premises prefix to AWS with the same BGP AS_PATH attribute across all the Site- to-Site VPN connections.

D. Advertise the on-premises prefix to AWS with a different BGP AS_PATH attribute across all the Site-to-Site VPN connections

E. Verify that the transit gateway attachments are present in the Availability Zones of the test VPC.

F. Verify that the on-premises location is sending traffic by using multiple flows.

Answer: A, B, F

Explanation:

Question: 112

A company is building a new workload on AWS that uses an Application Load Balancer (ALB) The company has configured a new ALB target group that uses slow start mode. A team begins registering Amazon EC2 Instances as targets in the new target group. During testing, the team observes that the targets did not enter slow start mode.

What caused the targets to not enter slow start mode?

A. The ALB configuration uses the round robin routing algorithm for traffic.

B. The target group did not contain at least one healthy target configured in slow start mode.

C. The target group must contain EC2 instances that are all the same instance type.

D. The ALB configuration uses the 5-tuple criteria for traffic.

Answer: B

Explanation:

Question: 113

A company is using third-party firewall appliances to monitor and inspect traffic on premises The company wants to use this same model on AWS. The company has a single VPC with an internet gateway. The VPC has a fleet of web servers that run on Amazon EC2 instances that are managed by an Auto Scaling group.

The company's network team needs to work with the security team to establish inline inspection of all packets that are sent to and from the web servers. The solution must scale as the fleet of virtual firewall appliances scales.

Which combination of steps should the network team take to implement this solution? (Select THREE.)

A. Create a new VPC, and deploy a fleet of firewall appliances. Create a Gateway Load Balancer. Add the firewall appliances as targets.

B. Create a security group for use with the firewall appliances, and allow port 443. Allow a port for the Gateway Load Balancer to perform health checks.

C. Create a security group for use with the firewall appliances, and allow port 6081. Allow a port for the Gateway Load Balancer to perform health checks.

D. Deploy a fleet of firewall appliances to the existing VPC. Create a Gateway Load Balancer. Add the firewall appliances as targets.

E. Update the internet gateway route table and the web server route table to send traffic to and from the internet to the VPC endpoint ID of the Gateway Load Balancer. Update the subnet route table that is associated with the Gateway Load Balancer endpoint to direct internet traffic to the internet gateway.

F. Create a new route table inside the web server VPC. Create a new edge association with the internet gateway. Update the internet gateway route table and the web server route table to send traffic to and from the internet to the VPC endpoint ID of the Gateway Load Balancer. Update the subnet route table that is associated with the Gateway Load Balancer endpoint to direct internet traffic to the internet gateway.

Answer: A, D, E

Explanation:

Question: 114

A company has a 2 Gbps AWS Direct Connect hosted connection from the company's office to a VPC in the ap-southeast-2 Region. A network engineer adds a 5 Gbps Direct Connect hosted connection from a different Direct Connect location in the same Region. The hosted connections are connected to different routers from the office with an iBGP session running in between the routers.

The network engineer wants to ensure that the VPC uses the 5 Gbps hosted connection to route traffic to the office. Failover to the 2 Gbps hosted connection must occur when the 5 Gbps hosted connection is down.

Which solution will meet these requirements?

A. Configure an outbound BGP policy from the router that is connected to the 2 Gbps connection. Advertise routes with a longer AS_PATH attribute to AWS.

B. Advertise a longer prefix route from the router that is connected to the 2 Gbps connection.

C. Advertise a less specific route from the router that is connected to the 5 Gbps connection.

D. Configure an outbound BGP policy from the router that is connected to the 5 Gbps connection. Advertise routes with a longer AS_PATH attribute to AWS.

Answer: A

Explanation:

Question: 115

A company has business operations in the United States and in Europe. The company's public applications are running on AWS and use three transit gateways. The transit gateways are located in the us-west-2. us-east-1. and eu-central-1 Regions. All the transit gateways are connected to each other in a full mesh configuration.

The company accidentally removes the route to the eu-central-1 VPCs from the us-west-2 transit gateway route table. The company also accidentally removes the route to the us-west-2 VPCs from the eu-central-1 transit gateway route table.

How can a network engineer identify the misconfiguration with the LEAST operational overhead?

A. Use the Route Analyzer feature for AWS Transit Gateway Network Manager

B. Use the AWSSupport-SetuplPMonitoringFromVPC AWS Systems Manager Automation runbook. Push network telemetry data to Amazon CloudWatch Logs for analysis.

C. Use VPC flow togs in eu-central-1 and us-west-2 to analyze the missing routes.

D. Use Amazon VPC Traffic Mirroring in eu-central-1 or us-west-2 to take packet captures and troubleshoot the connectivity issues.

Answer: C

Explanation:

Question: 116

Two companies are merging. The companies have a large AWS presence with multiple VPCs and are designing connectivity between their AWS networks. Both companies are using AWS Direct Connect with a Direct Connect gateway. Each company also has a transit gateway and multiple AWS Site-to- Site VPN connections from its transit gateway to on-premises resources. The new solution must optimize network visibility, throughput, logging, and monitoring.

Which solution will meet these requirements?

A. Configure a Site-to-Site VPN connection between each company's transit gateway to establish reachability between the respective networks. Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use VPC Reachability Analyzer to monitor connectivity.

B. Configure a Site-to-Site VPN connection between each company's transit gateway to establish reachability between the respective networks. Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use AWS Transit Gateway Network Manager to monitor the transit gateways and their respective connections.

C. Configure transit gateway peering between each company's transit gateway Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use VPC Reachability Analyzer to monitor connectivity.

D. Configure transit gateway peering between each company's transit gateway. Configure VPC Flow Logs for all VPCs. Publish the flow logs to Amazon CloudWatch. Use AWS Transit Gateway Network Manager to monitor the transit gateways, their respective connections, and the transit gateway peering link.

Answer: D

Explanation:

Question: 117

A company has a VPC in the AWS Cloud. The company recently acquired a competitor that also has a VPC in the AWS Cloud. A network engineer discovers an IP address overlap between the two VPCs. Both VPCs require access to an AWS Marketplace partner service.

Which solution will ensure interoperability among the VPC hosted services and the AWS Marketplace partner service?

A. Configure VPC peering with static routing between the VPCs. Configure an AWS Site-to-Site VPN connection with static routing to the partner service.

B. Configure a NAT gateway in the VPCs. Configure default routes in each VPC to point to the local NAT gateway. Attach each NAT gateway to a transit gateway. Configure an AWS Site-to-Site VPN connection with static routing to the partner service.

C. Configure AWS PrivateLink to facilitate connectivity between the VPCs and the partner service. Use the DNS name that is created with the associated interface endpoints to route traffic between the VPCs and the partner service.

D. Configure a NAT instance in the VPCs. Configure default routes in each VPC to point to the local NAT instance. Configure an interface endpoint in each VPC to connect to the partner service. Use the DNS name that is created with the associated interface endpoints to route traffic between the VPCs and the partner service.

Answer: C

Explanation:

Question: 118

A company has developed a new web application on AWS. The application runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate behind an Application Load Balancer (ALB) in the us- east-1 Region. The application uses Amazon Route 53 to host the DNS records for the domain. The content that is served from the website is mostly static images and files that are not updated frequently. Most of the traffic to the website from end users will originate from the United States. Some traffic will originate from Canada and Europe.

A network engineer needs to design a solution that will reduce latency for end users at the lowest cost. The solution also must ensure that all traffic is encrypted in transit until the traffic reaches the ALB.

Which solution will meet these requirements?

A. Configure the ALB to use an AWS Global Accelerator accelerator In us-east-1. Create a secure HTTPS listener. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to route to the DNS name that is assigned to the accelerator for the ALB.

B. Configure the ALB to use a secure HTTPS listener Create an Amazon CloudFront distribution. Set the origin domain name to point to the DNS record that is assigned to the ALB. Configure the CloudFront distribution to use an SSL certificate. Set all behaviors to force HTTPS. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to route to the DNS name that is assigned to the ALB.

C. Configure the ALB to use a secure HTTPS listener. Create an Amazon CloudFront distribution. Set the origin domain name to point to the DNS record that is assigned to the ALB. Configure the CloudFront distribution to use an SSL certificate and redirect HTTP to HTTPS. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to route to the CloudFront distribution.

D. Configure the ALB to use an AWS Global Accelerator accelerator in us-east-1. Create a secure HTTPS listener. Create a second application stack on Amazon ECS on Fargate in the eu-west-1 Region Create another secure HTTPS listener. Create an alias record in Amazon Route 53 for the custom domain name. Configure the alias record to use a latency-based routing policy to route to the DNS name that is assigned to the accelerator for the ALBs.

Answer: C

Explanation:

Question: 119

A financial company offers investment forecasts and recommendations to authorized users through the internet. All the services are hosted in the AWS Cloud. A new compliance requirement states that all the internet service traffic from any host must be logged and retained for 2 years. In its development AWS accounts, the company has designed, tested, and verified a solution that uses Amazon VPC Traffic Mirroring with a Network Load Balancer (NLB) as the traffic mirror target. While the solution runs in one AWS account, the solution mirrors the traffic to another AWS account.

A network engineer notices that not all traffic is mirrored when the solution is deployed into the production environment. The

network engineer also notices that this behavior is random.

Which statements are possible explanations for why not all the traffic is mirrored? (Select TWO.)

A. The security groups are misconflgured on the production AWS account that hosts the company's services.

B. The Amazon EC2 instance that is being monitored cannot handle the extra traffic that Traffic Mirroring has introduced

C. The 1AM policy that allows the creation of traffic mirror sessions is misconfigured.

D. The mirrored traffic has a lower priority than the production traffic and is being dropped when network congestion occurs.

E. The NLB is experiencing warm-up delay because of sudden and significant increases in traffic.

Answer: C, E

Explanation:

Question: 120

A company uses Amazon Route 53 to host a public hosted zone for example.com. A network engineer recently reduced the TTL on several records to 60 seconds. The network engineer wants to assess whether the change has increased the number of queries to Route 53 beyond the expected levels that the company identified before the change. The network engineer must obtain the number of queries that have been made to the example.com public hosted zone.

Which solution will provide this information?

A. Create a new trail in AWS CloudTrail to include Route 53 data events. Send logs to Amazon CloudWatch Logs. Set up a CloudWatch metric filter to count the number of queries and create graphs.

B. Use Amazon CloudWatch to access the AWS/Route 53 namespace and to check the DNSQuenes metric tor the public hosted zone.

C. Use Amazon CloudWatch to access the AWS/Route 53 Resolver namespace and to check the InboundQueryVolume metric for a specific endpoint.

D. Configure logging to Amazon CloudWatch for the public hosted zone. Set up a CloudWatch metric filter to count the number of queries and create graphs.

Answer: B

Explanation:

Question: 121

A banking company has an application that must connect to specific public IP addresses from a VPC. A network engineer has configured routes in the route table that is associated with the application's subnet to the required public IP addresses through an internet gateway.

The network engineer needs to set up email notifications that will alert the network engineer when a user adds a default route to the application subnet's route table with the internet gateway as a target.

Which solution will meet these requirements with the LEAST implementation effort?

A. Create an AWS Lambda function that reads the routes in the route table and sends an email notification. Configure the Lambda function to send an email notification if any route is configured with 0.0.0.0/0 or ::/0 CIDRs to the internet gateway. Configure the Lambda function to run every minute.

B. Create an AWS Lambda function that will be invoked by an Amazon EC2 CreateRoute API call. Configure the Lambda function to send an email notification. Configure the Lambda function to send an email notification if any route is configured with 0.0.0.0/0 or ::/0 CIDRs to the internet gateway.

C. Create AWS Config rules for the route table by using the internet-gateway-authorized-vpc-only managed rule. Create an Amazon EventBridge rule to match the AWS Config rule and to route to an Amazon Simple Notification Service (Amazon SNS) topic to send an email notification.

D. Create an AWS Config rule for the route table by using the no-unrestricted-route-to-igw managed rule. Create an Amazon EventBridge rule to match the AWS Config rule and to route to an Amazon Simple Notification Service (Amazon SNS) topic to send an email notification.

Answer: C

Explanation:

Question: 122

An ecommerce company needs to Implement additional security controls on all its domain names that are hosted in Amazon Route 53. The company's new policy requires data authentication and data integrity verification for all queries to the company's domain names. The current Route 53 architecture has four public hosted zones.

A network engineer needs to implement DNS Security Extensions (DNSSEC) signing and validation on the hosted zones. The solution must include an alert capability.

Which combination of steps will meet these requirements? {Select THREE.)

A. Enable DNSSEC signing for Route 53. Request that Route 53 create a Key-signing key (KSK) based on a customer managed key in AWS Key Management Service (AWS KMS).

B. Enable DNSSEC signing for Route 53. Request that Route 53 create a zone-signing key (ZSK) based on a customer managed key in AWS Key Management Service (AWS KMS).

C. Create a chain of trust for the hosted zones by adding a Delegation Signer (DS) record for each subdomain.

D. Create a chain of trust for the hosted zones by adding a Delegation Signer (DS) record to the parent zone.

E. Set up an Amazon CloudWatch alarm that provides an alert whenever a DNSSECInternalFailure error or DNSSECKeySigningKeysNeedingAction error is detected.

F. Set up an AWS CloudTrail alarm that provides an alert whenever a DNSSECInternalFailure error or DNSSECKeySigningKeysNeedingAction error is detected.

Answer: A, D, E

Explanation:

Question: 123

A company uses the us-east-1 Region and the ap-south-1 Region for its business units (BUs). The BUs are named BU-1 and BU-2. For each BU. there are two VPCs in us-east-1 and one VPC in ap-south-1.

Because of workload isolation requirements, resources can communicate within the same BU but cannot communicate with resources in the other BU. The company plans to add more BUs and plans to expand into more Regions.

Which solution will meet these requirements with the MOST operational efficiency?

A. Configure an AWS Cloud WAN network that operates in the required Regions Attach all BU VPCs to the AWS Cloud WAN core network. Update the AWS Cloud WAN segment actions to configure new routes to deny traffic between the different BU segments.

B. Configure a transit gateway in each Region. Configure peering between the transit gateways. Attach the BU VPCs to the transit gateway in the corresponding Region. Configure the transit gateway and VPC route tables to isolate traffic between BU VPCs.

C. Configure an AWS Cloud WAN network that operates in the required Regions. Attach all BU VPCs to the AWS Cloud WAN core network. Update the core network policy by setting the isolate- attachments parameter for each segment.

D. Configure an AWS Cloud WAN network that operates in the required Regions. Create AWS Cloud WAN segments for each BU. Configure VPC attachments for each BU's VPCs to the corresponding BU segment.

Answer: D

Explanation:

Question: 124

A company's network engineer must implement a cloud-based networking environment for a network operations team to centrally manage. Other teams will use the environment. Each team must be able to deploy infrastructure to the environment and must be able to manage its own resources. The environment must feature IPv4 and IPv6 support and must provide internet connectivity in a dual-stack configuration.

The company has an organization in AWS Organizations that contains a workload account for the teams. The network engineer creates a new networking account in the organization.

Which combination of steps should the network engineer take next to meet the requirements?

(Select THREE.)

A. Create a new VPC. Associate an IPv4 CIDR block of 10.0.0.0/16 and specify an IPv6 block of 2001: db8:c5a:6000::/56. Provision subnets by assigning /24 IPv4 CIDR blocks and /64 IPv6 CIDR blocks.

B. Create a new VPC. Associate an IPv4 CIDR block of 10.0.0.0/16 and use an Amazon-provided IPv6 CIDR block. Provision subnets by assigning /24 IPv4 CIDR blocks and 164 IPv6 CIDR blocks.

C. Enable sharing of resources within the organization by using AWS Resource Access Manager (AWS RAM). Create a resource share in the networking account, select the provisioned subnets, and share the provisioned subnets with the target workload account. Use the workload account to accept the resource share through AWS RAM.

D. Enable sharing of resources within the organization by using AWS Resource Access Manager (AWS RAM). Create a resource share in the networking account, select the new VPC. and share the new VPC with the target workload account. Use the workload account to accept the resource share through AWS RAM.

E. Create an internet gateway and an egress-only internet gateway. Deploy NAT gateways to the public subnets. Associate the internet gateway with the new VPC. Update the route tables. Associate the route tables with the relevant subnets.

F. Create an internet gateway. Deploy NAT instances to public subnets. Update the route tables. Associate the route tables with the relevant subnets.

Answer: A, C, E

Explanation:

Question: 125

A company's AWS infrastructure is spread across more than 50 accounts and across five AWS Regions. The company needs to manage its security posture with simplified administration and maintenance for all the AWS accounts. The company wants to use AWS Firewall Manager to manage the firewall rules and requirements.

The company creates an organization with all features enabled in AWS Organizations.

Which combination of steps should the company take next to meet the requirements? (Select THREE.)

A. Configure only the Firewall Manager administrator account to join the organization.

B. Configure all the accounts to join the organization.

C. Set an account as the Firewall Manager administrator account.

D. Set an account as the Firewall Manager child account.

E. Set up AWS Config for all the accounts and all the Regions where the company has resources.

F. Set up AWS Config for only the organization's management account.

Answer: B, C, E

Explanation:

Question: 126

A network engineer is working on a private DNS design to integrate AWS workloads and on-premises resources. The AWS deployment consists of five VPCs in the eu-west-1 Region that connect to the onpremises network over AWS Direct Connect. The VPCs communicate with each other by using a transit gateway. Each VPC is associated with a private hosted zone that uses the aws.example.internal domain. The network engineer creates an Amazon Route 53 Resolver outbound endpoint in a shared services VPC and attaches the shared services VPC to the transit gateway.

The network engineer is implementing a solution for DNS resolution. Queries for hostnames that end with aws.example.internal must use the private hosted zone. Queries for hostnames that end with all other domains must be forwarded to a private on-premises DNS resolver.

Which solution will meet these requirements?

A. Add a forwarding rule for”””that targets the on-premises server's DNS IP address. Add a system rule for aws.example.internal that targets Route 53 Resolver.

B. Add a forwarding rule for aws example.internal that targets Route 53 Resolver. Add a system rule for V that targets the Route 53 Resolver outbound endpoint.

C. Add a forwarding rule for”””that targets the Route 53 Resolver outbound endpoint.

D. Add a forwarding rule for"." that targets the Route 53 Resolver outbound endpoint.

Answer: D

Explanation:

Question: 127

A company is migrating its on-premises network from its data center in Virginia to its data center in New York. The AWS Direct

Connect connections for the Virginia and New York data center locations are both associated to the us-east-1 Region. The company needs to migrate a private VIF on an existing Direct Connect hosted connection from Virginia to New York. The company's on-premises network uses the connection to access VPCs through a Direct Connect gateway in us-east-1.

The company has already requested a new Direct Connect hosted connection from the new data center to the New York Direct Connect location.

Which solution will meet these requirements with the LEAST downtime?

A. Create a new private VIF on the new Direct Connect hosted connection Create a new Direct Connect gateway and attach the gateway to the new private VIF. Configure BGP routing on the new private VIF as a backup route. Perform the switchover during a maintenance window by shutting down BGP on the existing private VIF. Decommission the existing Direct Connect connection.

B. Create a new private VIF on the new Direct Connect hosted connection. Attach the new private VIF to the existing Direct Connect gateway. Configure BGP routing on the new private VIF as a backup route. Perform the switchover during a maintenance window by shutting down BGP on the existing private VIF. Decommission the existing Direct Connect connection.

C. During a maintenance window, migrate the existing private VIF to the new Direct Connect hosted connection. Attach the existing private VIF to the existing Direct Connect gateway. Decommission the existing Direct Connect connection.

D. During a maintenance window, delete the existing private VIF and create a new private VIF to the new Direct Connect hosted connection. Attach the new private VIF to the existing Direct Connect gateway. Decommission the existing Direct Connect hosted connection.

Answer: B

Explanation:

Question: 128

A company runs an application on Amazon EC2 instances. A network engineer implements a NAT gateway in the application's VPC to replace self-managed NAT instances. After the network engineer shifts traffic from the self-managed NAT instances to the NAT gateway, users begin to report issues.

During troubleshooting, the network engineer discovers that the connection to the application is closing after approximately 6 minutes of inactivity.

What should the network engineer do to resolve this issue?

A. Check for increases in the Amazon CloudWatch IdleTimeoutCount metric for the NAT gateway.
Configure TCP keepalive on the application EC2 instances.

B. Check for increases in the Amazon CloudWatch ErrorPortAllocation metric for the NAT gateway.
Configure an HTTP timeout value on the application EC2 instances.

C. Check for increases in the Amazon CloudWatch PacketsDropCount metric for the NAT gateway.

Configure an HTTPS timeout value on the application EC2 instances.

D. Check for decreases in the Amazon CloudWatch ActiveConnectionCount metric for the NAT gateway. Configure UDP keepalive on the application EC2 instances.

Answer: A

Explanation:

Question: 129

A network engineer needs to improve the network security of an existing AWS environment by adding an AWS Network Firewall firewall to control internet-bound traffic. The AWS environment consists of five VPCs. Each VPC has an internet gateway. NAT gateways, public Application Load Balancers (ALBs), and Amazon EC2 instances. The EC2 instances are deployed in private subnets. The architecture is deployed across two Availability Zones.

The network engineer must be able to configure rules for the public IP addresses in the environment, regardless of the direction of traffic. The network engineer must add the firewall by implementing a solution that minimizes changes to the existing production environment. The solution also must ensure high availability.

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

A. Create a centralized inspection VPC with subnets in two Availability Zones. Deploy Network Firewall in this inspection VPC with an endpoint in each Availability Zone.

B. Configure new subnets in two Availability Zones in each VPC. Deploy Network Firewall in each VPC with an endpoint in each Availability Zone.

C. Deploy Network Firewall in each VPC. Use existing subnets in each of the two Availability Zones to deploy Network Firewall endpoints.

D. Update the route tables that are associated with the private subnets that host the EC2 instances. Add routes to the Network Firewall endpoints.

E. Update the route tables that are associated with the public subnets that host the NAT gateways and the ALBs. Add routes to the Network Firewall endpoints.

Answer: C, D

Explanation:

Question: 130

A company has deployed a multi-VPC environment in the AWS Cloud. The company uses a transit gateway to connect all the VPCs together. In the past, the company has experienced a loss of connectivity between applications after changes to security groups, network ACLs. and route tables in a VPC. When these changes occur, the company wants to automatically verify that connectivity still exists between different resources in a single VPC.

Which solution will meet these requirements?

A. Create a list of paths between different resources to check in VPC Reachability Analyzer. Create an Amazon EventBridge rule to monitor when a change is made and logged in Amazon CloudWatch. Configure the rule to invoke an AWS Lambda function to test the different paths in Reachability Analyzer.

B. Create a list of paths between different resources to check in VPC Reachability Analyzer. Create an Amazon EventBridge rule to monitor when a change is made and logged in AWS CloudTrail. Configure the rule to invoke an AWS Lambda function to test the different paths in Reachability Analyzer.

C. Create a list of paths to check in AWS Network Manager Route Analyzer. Create an Amazon EventBridge rule to monitor when a change is made and logged in Amazon CloudWatch. Configure the rule to invoke an AWS Lambda function to test the different paths in Route Analyzer.

D. Create a list of paths to check in AWS Network Manager Route Analyzer. Create an Amazon EventBridge rule to monitor when a change is made and logged in AWS CloudTrail. Configure the rule to invoke an AWS Lambda function to test the different paths in Route Analyzer.

Answer: B

Explanation:

Question: 131

A company is planning to host external websites on AWS. The websites will include multiple tiers such as web servers, application logic services, and databases. The company wants to use AWS Network Firewall. AWS WAR and VPC security groups for network security.

The company must ensure that the Network Firewall firewalls are deployed appropriately within relevant VPCs. The company needs the ability to centrally manage policies that are deployed to Network Firewall and AWS WAF rules. The company also needs to allow application teams to manage their own security groups while ensuring that the security groups do not allow overly permissive access.

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

A. Define Network Firewall firewalls. AWS WAFv2 web ACLs. Network Firewall policies, and VPC security groups in code Use AWS CloudFormation to deploy the objects and Initial policies and rule groups. Use CloudFormation to update the AWS WAFv2 web ACLs. Network Firewall policies, and VPC security groups. Use Amazon GuardDuty to monitor for overly permissive rules.

B. Define Network Firewall firewalls. AWS WAFv2 web ACLs. Network Firewall policies, and VPC security groups in code. Use the AWS Management Console or the AWS CLI to manage the AWS WAFv2 web ACLs. Network Firewall policies, and VPC security

groups. Use Amazon GuardDuty to invoke an AWS Lambda function to evaluate the configured rules and remove any overly permissive rules.

C. Deploy AWS WAFv2 IP sets and AWS WAFv2 web ACLs with AWS CloudFormation. Use AWS Firewall Manager to deploy Network Firewall firewalls and VPC security groups where required and to manage the AWS WAFv2 web ACLs, Network Firewall policies, and VPC security groups.

D. Define Network Firewall firewalls. AWS WAFv2 web ACLs. Network Firewall policies, and VPC security groups in code. Use AWS CloudFormation to deploy the objects and initial policies and rule groups. Use AWS Firewall Manager to manage the AWS WAFv2 web ACLs. Network Firewall policies, and VPC security groups. Use Amazon GuardDuty to monitor for overly permissive rules.

Answer: D

Explanation:

Question: 132

A company is developing an API-based application on AWS for its process workflow requirements. The API will be invoked by clients in the company's on-premises data centers. The company has set up an AWS Direct Connect connection between on premises and AWS. A network engineer decides to implement the API as a private REST API in Amazon API Gateway. The network engineer wants to ensure that clients can reach the API endpoint through private communication.

Which solution can the network engineer use to invoke the API without any additional infrastructure setup?

A. Create an interface VPC endpoint for API Gateway with private DNS names enabled. Access the API by using the private DNS name of the endpoint.

B. Create an interface VPC endpoint for API Gateway with private DNS names enabled. Access the API by using an Amazon Route 53 alias of the endpoint.

C. Create an interface VPC endpoint for API Gateway. Associate the endpoint with the private REST API. Access the API by using an Amazon Route 53 alias of the endpoint.

D. Create an interface VPC endpoint for API Gateway with private DNS names enabled. Access the API by using the public DNS name of the endpoint.

Answer: A

Explanation:

Question: 133

A company has an internal web-based application that employees use. The company hosts the application over a VPN in the company's on-premises network. The application runs on a fleet of Amazon EC2 instances in a private subnet behind a Network Load Balancer (NLB) in the same subnet. The instances are in an Amazon EC2 Auto Scaling group.

During a recent security incident, SQL injection occurred on the application. A network engineer must implement a solution to prevent SQL injection attacks in the future.

Which combination of steps will meet these requirements? (Select THREE.)

A. Create an AWS WAF web ACL that includes rules to block SQL injection attacks

B. Create an Amazon CloudFront distribution. Specify the EC2 instances as the origin.

C. Replace the NLB with an Application Load Balancer

D. Associate the AWS WAF web ACL with the NLB.

E. Associate the AWS WAF web ACL with the Application Load Balancer.

F. Associate the AWS WAF web ACL with the Amazon CloudFront distribution.

Answer: A, D, E

Explanation:

Question: 134

A retail company is migrating its on-premises application to the AWS Cloud. Currently, the company has two on-premises data center locations. One data center is on the east coast of the United States, and one data center is on the west coast.

Each data center hosts four database systems. The largest database system stores 500 GB of data. The data centers are interconnected by two 10 GbE circuits for data synchronization. Each data center has two separate 1 GbE upstream internet connections. The company plans to have eight total VPCs to service its multiple business units. Four VPCs will be in the us-east-1 Region, and four will be in the us-west-2 Region.

A network engineer needs to design a connectivity solution that allows VPC-to-VPC connectivity. The solution must also allow secure connections between the on-premises data centers and AWS during the migration process. The company expects spikes in traffic among the VPCs during database synchronization. The company wants to run the migration plan during one weekend and as soon as technically possible. The company also wants to minimize long-term operational and human resources costs.

Which combination of steps will meet these requirements? {Select TWO.)

A. Deploy one transit gateway and attach all VPCs to it. Update the transit gateway and VPC route tables to allow any VPC to connect to any other VPC.

B. Configure VPC peering between all the VPCs. Update the VPC route tables to allow connectivity.

C. Provision two AWS Direct Connect connections from two Direct Connect locations that serve us- east-1 and us-west-2 to

provide connectivity between the data centers and AWS.

D. Provision one transit gateway VPN attachment for each data center to build connectivity between the on-premises data centers and AWS VPCs.

E. Provision one AWS Site-to-Site VPN connection for each data center and for each VPC to build connectivity between the on-premises data centers and AWS VPCs.

Answer: A, C

Explanation:

Question: 135

A company has a hybrid IT setup that includes services that run in an on-premises data center and in the AWS Cloud. The company is using AWS Direct Connect to connect its data center to AWS. The company is using one AWS Site-to-Site VPN connection as backup and requires a backup connectivity option to always be present. The company is transitioning to IPv6 by implementing dual-stack architectures.

Which combination of steps will transition the data center's connectivity to AWS in the LEAST amount of time? (Select TWO.)

A. Create a new Site-to-Site VPN tunnel for the IPv6 traffic.

B. Create a new dual-stack Site-to-Site VPN connection between the data center and AWS. Provision routing. Delete the original Site-to-Site VPN connection

C. Associate a new dual-stack public VIF with the Direct Connect connection. Migrate the Direct Connect traffic to the new VIF.

D. Add a new IPv6 peer in the existing VIF. Use the IPv6 address provided by Amazon on the peer router.

E. Send IPv6 traffic between the data center and AWS in a tunnel inside the existing IPv4 tunnels.

Answer: B, C

Explanation:

Question: 136

A network engineer needs to design the architecture for a high performance computing (HPC) workload. Amazon EC2 instances will require 10 Gbps flows and an aggregate throughput of up to 100 Gbps across many instances with low-latency communication.

Which architecture solution will optimize this workload?

A. Place nodes in a single subnet of a VPC. Configure a cluster placement group. Ensure that the latest Elastic Fabric Adapter (EFA) drivers are installed on the EC2 instances with a supported operating system.

B. Place nodes in multiple subnets in a single VPC. Configure a spread placement group Ensure that the EC2 instances support Elastic Network Adapters (ENAs) and that the drivers are updated on each instance operating system.

C. Place nodes in multiple VPCs. Use AWS Transit Gateway to route traffic between the VPCs. Ensure that the latest Elastic Fabric Adapter (EFA) drivers are installed on the EC2 instances with a supported operating system.

D. Place nodes in multiple subnets in multiple Availability Zones. Configure a cluster placement group. Ensure that the EC2 instances support Elastic Network Adapters (ENAs) and that the drivers are updated on each instance operating system.

Answer: A

Explanation:

Question: 137

An online retail company is running a web application in the us-west-2 Region and serves consumers in the United States. The company plans to expand across several countries in Europe and wants to provide low latency for all its users.

The application needs to identify the users' IP addresses and provide localized content based on the users' geographic location. The application uses HTTP GET and POST methods for its functionality.

The company also needs to develop a failover mechanism that works for GET and POST methods and is based on health checks. The failover must occur in less than 1 minute for all clients.

Which solution will meet these requirements?

A. Configure a Network Load Balancer (NLB) for the application in each environment In the new AWS Regions. Create an AWS Global Accelerator accelerator that has endpoint groups that point to the NLBs in each Region.

B. Configure an Application Load Balancer (ALB) for the application in each environment in the new AWS Regions. Create an AWS Global Accelerator accelerator that has endpoint groups that point to the ALBs in each Region.

C. Configure an Application Load Balancer (ALB) for the application in each environment in the new AWS Regions. Create Amazon Route 53 public hosted zones that have failover routing policies.

D. Configure a Network Load Balancer (NLB) for the application in each environment in the new AWS Regions. Create an Amazon CloudFront distribution. Configure an origin group with origin failover options.

Answer: B

Explanation:

Question: 138

A company has a VPC that hosts Amazon EC2 instances in a private subnet. The EC2 Instances use a NAT gateway and an internet gateway for internet connectivity to retrieve data from specific internet websites. The company wants to use AWS Network Firewall to filter outbound traffic.

What should a network engineer do to meet these requirements?

A. 1. Create a firewall in the NAT gateway subnet.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.

3. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.

4. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

B. 1. Create a firewall in a new subnet.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.

3. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.

4. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

C. 1. Create a firewall in the subnet of the EC2 instances.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.

3. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.

4. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

D. 1. Create a firewall in a new subnet.

2. Configure the EC2 instance subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the NAT gateway.

3. Configure the NAT gateway subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the firewall endpoint.

4. Configure the firewall subnet route tables to direct traffic with a destination of 0.0.0.0/0 to the internet gateway.

Answer: B

Explanation:

Question: 139

A company is deploying a new stateless web application on AWS. The web application will run on Amazon EC2 instances in private subnets behind an Application Load Balancer. The EC2 instances are in an Auto Scaling group. The web application has a stateful management application for administration that will run on EC2 instances that are in a separate Auto Scaling group.

The company wants to access the management application by using the same URL as the web application, with a path prefix of

/management. The protocol, hostname, and port number must be the same for the web application and the management application. Access to the management application must be restricted to the company's on-premises IP address space. An SSL/TLS certificate from AWS Certificate Manager (ACM) will protect the web application.

Which combination of steps should a network engineer take to meet these requirements? (Select TWO.)

A. Insert a rule for the load balancer HTTPS listener. Configure the rule to check the path-pattern condition type for the /management prefix and to check the source-ip condition type for the onpremises IP address space. Forward requests to the management application target group if there is a match. Edit the management application target group and enable stickiness.

B. Modify the default rule for the load balancer HTTPS listener. Configure the rule to check the pathpattern condition type for the /management prefix and to check the source-Ip condition type for the on-premises IP address space. Forward requests to the management application target group if there is not a match. Enable group-level stickiness in the rule attributes.

C. Insert a rule for the load balancer HTTPS listener. Configure the rule to check the path-pattern condition type for the /management prefix and to check the X-Forwarded-For HTTP header for the on-premises IP address space. Forward requests to the management application target group if there is a match. Enable group-level stickiness in the rule attributes.

D. Modify the default rule for the load balancer HTTPS listener. Configure the rule to check the pathpattern condition type for the /management prefix and to check the source-Ip condition type for the on-premises IP address space. Forward requests to the web application target group if there is not a match.

E. Forward all requests to the web application target group. Edit the web application target group and disable stickiness.

Answer: A, E

Explanation:

Question: 140

A company has a data center in the us-west-1 Region with a 10 Gbps AWS Direct Connect dedicated connection to a Direct Connect gateway. There are two private VIFs from the same data center location in us-west-1 that are attached to the same Direct Connect gateway.

VIF 1 advertises 172.16.0.0/16 with an AS PATH attribute value of 65000. VIF 2 advertises 172.16.1.0/24 with an AS PATH attribute value of 65000 65000 65000.

How will AWS route traffic to the data center for traffic that has a destination address within the 172.16.1.0/24 network range?

A. AWS will route all traffic by using VIF 1.

B. AWS will route all traffic by using VIF 2.

C. AWS will use both VIFs for routing by using a round-robin policy.

D. AWS will use flow control to balance the traffic between the two VIFs.

Answer: B

Explanation:

Question: 141

A development team is building a new web application in the AWS Cloud. The main company domain, example.com. is currently hosted in an Amazon Route 53 public hosted zone in one of the company's production AWS accounts.

The developers want to test the web application in the company's staging AWS account by using publicly resolvable subdomains under the example.com domain with the ability to create and delete DNS records as needed. Developers have full access to Route 53 hosted zones within the staging account, but they are prohibited from accessing resources in any of the production AWS accounts.

Which combination of steps should a network engineer take to allow the developers to create records under the example.com domain? (Select TWO.)

A. Create a public hosted zone for example.com in the staging account.

B. Create a staging.example.com NS record in the example.com domain. Populate the value with the name servers from the staging.example.com domain. Set the routing policy type to simple routing.

C. Create a private hosted zone for stagmg.example.com in the staging account.

D. Create an example.com NS record in the staging.example.com domain. Populate the value with the name servers from the example.com domain. Set the routing policy type to simple routing

E. Create a public hosted zone for staging.example.com in the staging account.

Answer: B, E

Explanation:

When a client queries a DNS server for a domain name, the DNS server typically starts by looking for NS records to determine which name servers are authoritative for the domain. The DNS server then queries the authoritative name servers to obtain the information about the domain that the client requested. For example, suppose you own the domain example.com, but you want to delegate control of the subdomain sub.example.com to a different set of name servers. You would create NS records in the example.com zone file that point to the name servers for sub.example.com. This tells DNS servers that the name servers for sub.example.com are authoritative for that subdomain, and they should query those name servers for any requests related to sub.example.com.

Question: 142

A company is deploying AWS Cloud WAN with edge locations in the us-east-1 Region and the ap- southeast-2 Region. Individual AWS Cloud WAN segments are configured for the development environment, the production environment, and the shared services environment at each edge location. Many new VPCs will be deployed for the environments and will be configured as attachments to the AWS Cloud WAN core network.

The company's network team wants to ensure that VPC attachments are configured for the correct segment. The network team will tag the VPC attachments by using the Environment key with a value of the corresponding environment segment name. The segment for the production environment in us-east-1 must require acceptance for attachment requests. AH other attachment requests must not require acceptance.

Which solution will meet these requirements?

A. Create a rule with a number of 100 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the "or" value. Include conditions that require a tag:Environment value of Production or a Region value of us-east-1. Create a rule with a number of 200 that does not require acceptance to map any tag:Environment values to their respective segments.

B. Create a rule with a number of 100 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the "and" value. Include conditions that require a tag:Environment value of Production and a Region value of us-east-1. Create a rule with a number of 200 that does not require acceptance to map any tag:Environment values to their respective segments.

C. Create a rule with a number of 100 that does not require acceptance to map any tag:Environment values to their respective segments. Create a rule with a number of 200 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the "and" value. Include conditions that require a tag:Environment value of Production and a Region value of us-east- 1.

D. Create a rule with a number of 100 that does not require acceptance to map any tag:Environment values to their respective segments Create a rule with a number of 200 that requires acceptance for attachments to the production segment. In the rule, set the condition logic to the "or value. Include conditions that require a tag:Environment value of Production or a Region value of us-east-1.

Answer: B

Explanation:

https://docs.aws.amazon.com/network-manager/latest/cloudwan/cloudwan-policy- attachments.html

Question: 143

A company has two data centers that are interconnected with multiple redundant links from different suppliers. The company uses IP addresses that are within the 172.16.0.0/16 CIDR block. The company is running iBGP between the two data centers by using a

private Autonomous System Number (ASN)and IGP.

The company is moving toward a hybrid setup in which the company will initially use one VPC in the AWS Cloud. An AWS Direct Connect connection runs from the first data center to a Direct Connect gateway by using a private VIF On the connection, the company advertises a summarized route for the 172.16.0.0/16 network The company is planning to set up a second summarized route from the second data center to a different Direct Connect location.

The company needs to implement a solution to route traffic to and from AWS through the first Direct Connect connection. The solution must use the second Direct Connect connection for failover purposes only.

Which solution will meet these requirements?

A. Prepend the private ASN on the BGP announcements to AWS from the second data center. Add a second VIF in the first Direct Connect connection. Advertise the same network without any prepends from the first data center. Implement the same setup for the BGP announcement from AWS to the two data centers.

B. Tag the BGP announcements with the local preference BGP community tags. Set the tag to high preference for the first data center. Set the tag to low preference for the second data center. Configure the second data center's router to have a lower local preference for the direct AWS BGP advertisements than for the advertisement from the first data center.

C. Configure the Direct Connect gateway to prefer routing through the Direct Connect connection with the first data center. Configure the second data center's router to have a lower local preference for the direct AWS BGP advertisements than for the advertisement from the first data center.

D. Configure the local AWS Region BGP community tag on the BGP route that is advertised from the first data center. Configure AS PATH prepends on the BGP announcements from the second data center.

Answer: C

Explanation:

Question: 144

A company hosts its IT infrastructure in an on-premises data center. The company wants to migrate the infrastructure to the AWS Cloud in phases. A network engineer wants to set up a 10 Gbps AWS Direct Connect dedicated connection between the on-premises data center and VPCs. The company's network provider needs 3 months to provision the Direct Connect connection.

In the meantime, the network engineer implements a temporary solution by deploying an AWS Site- to-Site VPN connection that terminates to a virtual private gateway. The network engineer observes that the bandwidth of the Site-to-Site VPN connection is capped at 1.25 Gbps despite a powerful customer gateway device.

What should the network engineer do to improve the VPN connection bandwidth before the implementation of the Direct Connect connection?

A. Contact AWS Support to request a bandwidth quota increase for the existing Site-to-Site VPN connection.

B. Discuss the issue with the hardware vendor. Buy a bigger and more powerful customer gateway device that has faster encryption and decryption capabilities.

C. Create several additional Site-to-Site VPN connections that terminate on the same virtual gateway. Configure equal-cost multi-path (ECMP) routing to use all the VPN connections simultaneously.

D. Create a transit gateway. Attach the VPCs to the transit gateway. Create several additional Site-to- Site VPN connections that terminate on the transit gateway. Configure equal-cost multi-path (ECMP) routing to use all the VPN connections simultaneously.

Answer: C

Explanation:

Question: 145

A company has an application that runs on premises. The application needs to communicate with an application that runs in a VPC on AWS. The communication between the applications must be encrypted and must use private IP addresses. The communication cannot travel across the public internet.

The company has established a 1 Gbps AWS Direct Connect connection between the on-premises location and AWS.

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

A. Configure a private VIF on the Direct Connect connection. Associate the private VIF with the VPC's virtual private gateway. Set up an AWS Site-to-Site VPN private IP VPN connection to the virtual private gateway.

B. Create a transit gateway. Configure a transit VIF on the Direct Connect connection. Associate the transit VIF with a Direct Connect gateway. Associate the Direct Connect gateway with a new transit gateway. Set up an AWS Site-to-Site VPN private IP VPN connection to the transit gateway.

C. Configure a public VIF on the Direct Connect connection. Associate the public VIF with a Direct Connect gateway. Associate the Direct Connect gateway with a new transit gateway. Set up an AWS Site-to-Site VPN private IP VPN connection to the transit gateway.

D. Create a transit gateway. Configure a transit VIF on the Direct Connect connection. Associate the transit VIF with a Direct Connect gateway. Associate the Direct Connect gateway with a new transit gateway. Set up a third-party firewall in a new VPC that is attached to the transit gateway. Set up a VPN connection to the third-party firewall.

Answer: D

Explanation:

Question: 146

A network engineer is using AWS Direct Connect connections and MACsec to encrypt data from a corporate data center to the

Direct Connect location. The network engineer learns that the MACsec secret key might have been compromised. The network engineer needs to update the connection with an uncompromised secure key.

Which solution will meet this requirement?

A. Create a new MACsec secret key that uses an AWS Key Management Service (AWS KMS) AWS managed key. Associate the new pre-shared key, Connection Key Name (CKN). and Connectivity Association Key (CAK) with the connection.

B. Create a new MACsec secret key that uses an AWS Key Management Service (AWS KMS) customer managed key. Associate the new pre-shared key, Connection Key Name (CKN). and Connectivity Association Key (CAK) with the connection.

C. Modify the existing MACsec secret key. Re-associate the existing pre-shared key. Connection Key Name (CKN), and Connectivity Association Key (CAK) with the connection.

D. Modify the existing MACsec secret key. Associate the new pre-shared key. Connection Key Name (CKN). and Connectivity Association Key (CAK) with the connection.

Answer: A

Explanation:

Question: 147

A company is running business applications on AWS. The company uses 50 AWS accounts, thousands of VPCs. and 3 AWS Regions across the United States and Europe.

A network engineer needs to establish network connectivity between an on-premises data center and the Regions. The network engineer also must establish connectivity between the VPCs. Onpremises users and applications must be able to connect to applications that run in the VPCs.

The company has an existing AWS Direct Connect connection that the network engineer can use. The network engineer creates a transit gateway in each Region and configures the transit gateways as inter-Region peers.

Which solution will provide network connectivity from the on-premises data center to the Regions and will provide inter-VPC communications across the different Regions?

A. Create a private VIF with a gateway type of virtual private gateway. Configure the private VIF to use a virtual private gateway that is associated with one of the VPCs.

B. Create a private VIF to a new Direct Connect gateway. Associate the new Direct Connect gateway with a virtual private gateway in each VPC.

C. Create a transit VIF with a gateway association to a new Direct Connect gateway. Associate each transit gateway with the new Direct Connect gateway.

D. Create an AWS Site-to-Site VPN connection that uses a public VIF for the Direct Connect connection Attach the Site-to-Site VPN

connection to the transit gateways.

Answer: C

Explanation:

Question: 148

A company needs to manage Amazon EC2 instances through command line interfaces for Linux hosts and Windows hosts. The EC2 instances are deployed in an environment in which there is no route to the internet. The company must implement role-based access control for management of the instances. The company has a standalone on-premises environment.

Which approach will meet these requirements with the LEAST maintenance overhead?

A. Set up an AWS Direct Connect connection between the on-premises environment and the VPC where the instances are deployed. Configure routing, security groups, and ACLs. Connect to the instances by using the Direct Connect connection.

B. Deploy and configure AWS Systems Manager Agent (SSM Agent) on each instance. Deploy VPC endpoints for Systems Manager Session Manager. Connect to the instances by using Session Manager.

C. Establish an AWS Site-to-Site VPN connection between the on-premises environment and the VPC where the instances are deployed. Configure routing, security groups, and ACLs. Connect to the instances by using the Site-to-Site VPN connection.

D. Deploy an appliance to the VPC where the instances are deployed. Assign a public IP address to the appliance. Configure security groups and ACLs. Connect to the instances by using the appliance as an intermediary.

Answer: B

Explanation:

Question: 149

A company has an application that runs on a fleet of Amazon EC2 instances. A new company regulation mandates that all network traffic to and from the EC2 instances must be sent to a centralized third-party EC2 appliance for content inspection.

Which solution will meet these requirements?

A. Configure VPC flow logs on each EC2 network Interface. Publish the flow logs to an Amazon S3 bucket. Create a third-party EC2 appliance to acquire flow logs from the S3 bucket. Log in to the appliance to monitor network content.

B. Create a third-party EC2 appliance in an Auto Scaling group fronted by a Network Load Balancer (NLB). Configure a mirror session. Specify the NLB as the mirror target. Specify a mirror filter to capture inbound and outbound traffic for the source of the mirror session, specify the EC2 elastic network interfaces for all the instances that host the application.

C. Configure a mirror session. Specify an Amazon Data Firehose delivery stream as the mirror target Specify a mirror filter to capture inbound and outbound traffic. For the source of the mirror session, specify the EC2 elastic network interfaces for all the instances that host the application Create a third-party EC2 appliance. Send all traffic to the appliance through the Firehose delivery stream for content inspection.

D. Configure VPC flow logs on each EC2 network interface. Send the logs to Amazon CloudWatch. Create a third-party EC2 appliance. Configure a CloudWatch filter to send the flow logs to Amazon Data Firehose to load the logs into the appliance.

Answer: D

Explanation:

Question: 150

A company has stateful security appliances that are deployed to multiple Availability Zones in a centralized shared services VPC. The AWS environment includes a transit gateway that is attached to application VPCs and the shared services VPC. The application VPCs have workloads that are deployed in private subnets across multiple Availability Zones. The stateful appliances in the shared services VPC inspect all east-west (VPC-to-VPC) traffic.

Users report that inter-VPC traffic to different Availability Zones is dropping. A network engineer verified this claim by issuing Internet Control Message Protocol (ICMP) pings between workloads in different Availability Zones across the application VPCs. The network engineer has ruled out security groups, stateful device configurations, and network ACLs as the cause of the dropped traffic.

What is causing the traffic to drop?

A. The stateful appliances and the transit gateway attachments are deployed in a separate subnet in the shared services VPC.

B. Appliance mode is not enabled on the transit gateway attachment to the shared services VPC

C. The stateful appliances and the transit gateway attachments are deployed in the same subnet in the shared services VPC.

D. Appliance mode is not enabled on the transit gateway attachment to the application VPCs.

Answer: B

Explanation:

Question: 151

A company is building an internet-facing application that is hosted on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.

The company is using the Amazon VPC Container Network Interface (CNI) plugin for Kubernetes for pod networking connectivity. The company needs to expose its application to the internet by using a Network Load Balancer (NLB). The pods that host the application must have visibility of the source IP address that is contained in the original packet that the NLB receives.

How should the network engineer configure the NLB and Amazon EKS settings to achieve these goals?

A. Specify the Ip target type for the NLB. Set the externalTrafficPolicy attribute to Local in the Kubernetes service specification.

B. Specify the instance target type for the NLB. Set the externalTrafficPolicy attribute to Cluster in the Kubernetes service specification

C. Specify the instance target type for the NLB. Set the externalTrafficPolicy attribute to Local in the Kubernetes service specification.

D. Specify the Ip target type for the NLB. Set the externalTrafficPolicy attribute to Cluster in the Kubernetes service specification

Answer: A

Explanation:

Question: 152

A company has an AWS account with four VPCs in the us-east-1 Region. The VPCs consist of a development VPC and three production VPCs that host various workloads.

The company has extended its on-premises data center to AWS with AWS Direct Connect by using a Direct Connect gateway. The company now wants to establish connectivity to its production VPCs and development VPC from on premises. The production VPCs are allowed to route data to each other. However, the development VPC must be isolated from the production VPCs. No data can flow between the development VPC and the production VPCs.

In preparation to implement this solution, a network engineer creates a transit gateway with a single transit gateway route table. Default route table association and default route table propagation are turned off. The network engineer attaches the production VPCs. the development VPC. and the Direct Connect gateway to the transit gateway. For each VPC route table, the network engineer adds a route to 0.0.0.0/0 with the transit gateway as the next destination.

Which combination of steps should the network engineer take next to complete this solution? (Select THREE.)

A. Associate the production VPC attachments with the existing transit gateway route table. Propagate the routes from these attachments.

B. Associate all the attachments with the existing transit gateway route table. Propagate the routes from these attachments.

C. Associate the Direct Connect gateway attachment with the existing transit gateway route table. Propagate the Direct Connect gateway attachment to this route table.

D. Change the security group inbound rules on the existing transit gateway network interfaces in the development VPC to allow connections to and from the on-premises CIDR range only.

E. Create a new transit gateway route table. Associate the new route table with the development VPC attachment. Propagate the Direct Connect gateway and development VPC attachment to the new route table.

F. Create a new transit gateway with default route table association and default route table propagation turned on. Attach the Direct Connect gateway and development VPC to the new transit gateway.

Answer: A, C, D

Explanation:

Question: 153

A company has AWS accounts in an organization in AWS Organizations. The company has implemented Amazon VPC IP Address Manager (IPAM)in its networking AWS account. The company is using AWS Resource Access Manager (AWS RAM) to share IPAM pools with other AWS accounts. The company has created a top-level pool with a CIDR block of 10.0.0.0/8. For each AWS account, the company has created an IPAM pool within the top-level pool.

A network engineer needs to implement a solution to ensure that users in each AWS account cannot create new VPCs. The solution also must prevent users from associating a CIDR block with existing VPCs unless the CIDR block is from the IPAM pool for that account.

Which solution will meet these requirements?

A. Create a new AWS Config rule to find all VPCs that are not configured to allocate their CIDR block from an IPAM pool. Invoke an AWS Lambda function to delete these VPCs.

B. Create a new SCP in Organizations. Add a condition that denies the CreateVpc and AssociateVpcCidrBlock Amazon EC2 actions if the Ipv4IpamPoolId context key value is not the ID of an IPAM pool.

C. Create an AWS Lambda function to check for and delete all VPCs that are not configured to allocate their CIDR block from an IPAM pool. Invoke the Lambda function at regular intervals.

D. Create an Amazon EventBridge rule to check for AWS CloudTrail events for the CreateVpc and AssociateVpcCidrBlock Amazon EC2 actions. Use the rule to invoke an AWS Lambda function to delete all VPCs that are not configured to allocate their CIDR block from an IPAM pool.

Answer: B

Explanation:

Question: 154

A company is creating new features for its ecommerce website. These features will use several microservices that are accessed through different paths. The microservices will run on Amazon Elastic Container Service (Amazon ECS). The company requires the use of HTTPS for all of its public websites. The application requires the customer's source IP addresses.

A network engineer must implement a load balancing strategy that meets these requirements.

Which combination of actions should the network engineer take to accomplish this goal? (Choose two.)

A. Use a Network Load Balancer

B. Retrieve client IP addresses by using the X-Forwarded-For header

C. Use AWS App Mesh load balancing

D. Retrieve client IP addresses by using the X-IP-Source header

E. Use an Application Load Balancer.

Answer: B, E

Explanation:

Question: 155

A company is migrating its containerized application to AWS. For the architecture the company will have an ingress VPC with a Network Load Balancer (NLB) to distribute the traffic to front-end pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The front end of the application will determine which user is requesting access and will send traffic to 1 of 10 services VPCs. Each services VPC will include an NLB that distributes traffic to the services pods in an EKS cluster.

The company is concerned about overall cost. User traffic will be responsible for more than 10 TB of data transfer from the ingress VPC to services VPCs every month. A network engineer needs to recommend how to design the communication between the VPCs.

Which solution will meet these requirements at the LOWEST cost?

A. Create a transit gateway. Peer each VPC to the transit gateway. Use zonal DNS names for the NLB in the services VPCs to minimize cross-AZ traffic from the ingress VPC to the services VPCs.

B. Create an AWS PrivateLink endpoint in every Availability Zone in the ingress VPC. Each PrivateLink endpoint will point to

the zonal DNS entry of the NLB in the services VPCs.

C. Create a VPC peering connection between the ingress VPC and each of the 10 services VPCs. Use zonal DNS names for the NLB in the services VPCs to minimize cross-AZ traffic from the ingress VPC to the services VPCs.

D. Create a transit gateway. Peer each VPC to the transit gateway. Turn off cross-AZ load balancing on the transit gateway. Use Regional DNS names for the NLB in the services VPCs.

Answer: C

Explanation:

Question: 156

A company has hundreds of Amazon EC2 instances that are running in two production VPCs across all Availability Zones in the us-east-1 Region. The production VPCs are named VPC A and VPC B.

A new security regulation requires all traffic between production VPCs to be inspected before the traffic is routed to its final destination. The company deploys a new shared VPC that contains a stateful firewall appliance and a transit gateway with a VPC attachment across all VPCs to route traffic between VPC A and VPC B through the firewall appliance for inspection. During testing, the company notices that the transit gateway is dropping the traffic whenever the traffic is between two Availability Zones.

What should a network engineer do to fix this issue with the LEAST management overhead?

A. In the shared VPC, replace the VPC attachment with a VPN attachment. Create a VPN tunnel between the transit gateway and the firewall appliance. Configure BGP.

B. Enable transit gateway appliance mode on the VPC attachment in VPC A and VPC B.

C. Enable transit gateway appliance mode on the VPC attachment in the shared VPC.

D. In the shared VPC, configure one VPC peering connection to VPC A and another VPC peering connection to VPC B.

Answer: C

Explanation:

Question: 157

A company needs to transfer data between its VPC and its on-premises data center. The data must travel through a connection that has dedicated bandwidth. The data also must be encrypted in transit. The company has been working with an AWS Partner Network (APN) Partner to establish the connection.

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

A. Request a hosted connection from the APN Partner.

B. Request a hosted public VIF from the APN Partner.

C. Create an AWS Site-to-Site VPN connection.

D. Create an AWS Client VPN connection.

E. Create a private VIF.

F. Create a public VIF.

Answer: A, C, F

Explanation:

Question: 158

A company's security guidelines state that all outbound traffic from a VPC to the company's onpremises data center must pass through a security appliance. The security appliance runs on an Amazon EC2 instance. A network engineer needs to improve the network performance between the on-premises data center and the security appliance.

Which actions should the network engineer take to meet these requirements? (Choose two.)

A. Use an EC2 instance that supports enhanced networking.

B. Send outbound traffic through a transit gateway.

C. Increase the EC2 instance size.

D. Place the EC2 instance in a placement group within the VPC.

E. Attach multiple elastic network interfaces to the EC2 instance.

Answer: A, C

Explanation:

Question: 159

A company's security guidelines state that all outbound traffic from a VPC to the company's onpremises data center must pass through a security appliance. The security appliance runs on an Amazon EC2 instance. A network engineer needs to improve the network performance between the on-premises data center and the security appliance.

Which actions should the network engineer take to meet these requirements? (Choose two.)

A. Use an EC2 instance that supports enhanced networking.

B. Send outbound traffic through a transit gateway.

C. Increase the EC2 instance size.

D. Place the EC2 instance in a placement qroup within the VPC.

E. Attach multiple elastic network interfaces to the EC2 instance.

Answer: A, C

Explanation:

Question: 160

A company's application team is unable to launch new resources into its VPC. A network engineer discovers that the VPC has run out of usable IP addresses. The VPC CIDR block is 172.16.0.0/16.

Which additional CIDR block can the network engineer attach to the VPC?

A. 172.17.0.0/29

B. 10.0.0.0/16

C. 172.17.0.0/16

D. 192.168.0.0/16

Answer: C

Explanation:

Question: 161

A financial trading company is using Amazon EC2 instances to run its trading platform. Part of the company's trading platform includes a third-party pricing service that the EC2 instances communicate with over UDP on port 50000.

Recently, the company has had problems with the pricing service. Some of the responses from the pricing service appear to be incorrectly formatted and are not being processed successfully. The third-party vendor requests access to the data that the pricing service is returning. The third-party vendor wants to capture request and response data for debugging by logging in to an EC2

instance that accesses the pricing service. The company prohibits direct access to production systems and requires all log analysis to be performed in a dedicated monitoring account.

Which set of steps should a network engineer take to capture the data and meet these requirements?

A.

1. Configure VPC flow logs to capture the data that flows in the VPC.
2. Send the data to an Amazon S3 bucket.
3. In the monitoring account, extract the data that flows to the EC2 instance's IP address and filter the traffic for the UDP data.
4. Provide the data to the third-party vendor.

B.

1. Configure a traffic mirror filter to capture the UDP data.
2. Configure Traffic Mirroring to capture the traffic for the EC2 instance's elastic network interface.
3. Configure a packet inspection package on a new EC2 instance in the production environment. Use the elastic network interface of the new EC2 instance as the target for the traffic mirror.
4. Extract the data by using the packet inspection package.
5. Provide the data to the third-party vendor.

C.

1. Configure a traffic mirror filter to capture the UDP data.
2. Configure Traffic Mirroring to capture the traffic for the EC2 instance's elastic network interface.
3. Configure a packet inspection package on a new EC2 instance in the monitoring account. Use the elastic network interface of the new EC2 instance as the target for the traffic mirror.
4. Extract the data by using the packet inspection package.
5. Provide the data to the third-party vendor.

D.

1. Create a new Amazon Elastic Block Store (Amazon EBS) volume. Attach the EBS volume to the EC2 instance.
2. Log in to the EC2 instance in the production environment. Run the tcpdump command to capture the UDP data on the EBS volume.
3. Export the data from the EBS volume to Amazon S3.

4. Provide the data to the third-party vendor.

Answer: C

Explanation:

Question: 162

A company's network engineer is configuring an AWS Site-to-Site VPN connection between a transit gateway and the company's on-premises network. The Site-to-Site VPN connection is configured to use BGP over two tunnels in active/active mode with equal-cost multi-path (ECMP) routing activated on the transit gateway.

When the network engineer attempts to send traffic from the on-premises network to an Amazon EC2 instance, traffic is sent over the first tunnel. However, return traffic is received over the second tunnel and is dropped at the customer gateway. The network engineer must resolve this issue without reducing the overall VPN bandwidth.

Which solution will meet these requirements?

A. Configure the customer gateway to use AS PATH prepending and local preference to prefer one tunnel over the other.

B. Configure the Site-to-Site VPN options to set the first tunnel as the primary tunnel to eliminate asymmetric routing.

C. Configure the virtual tunnel interfaces on the customer gateway to allow asymmetric routing.

D. Configure the Site-to-Site VPN to use static routing in active/active mode to ensure that traffic flows over a preferred path.

Answer: D

Explanation:

Question: 163

A software-as-a-service (SaaS) company is migrating its private SaaS application to AWS. The company has hundreds of customers that connect to multiple data centers by using VPN tunnels. As the number of customers has grown, the company has experienced more difficulty in its effort to manage routing and segmentation of customers with complex NAT rules.

After the migration to AWS is complete, the company's AWS customers must be able to access the SaaS application directly from their VPCs. Meanwhile, the company's on-premises customers still must be able to connect through IPsec encrypted tunnels.

Which solution will meet these requirements?

A. Connect the AWS customer VPCs to a shared transit gateway. Use AWS Site-to-Site VPN connections to the transit gateway for the on-premises customers

B. Use AWS PrivateLink to connect the AWS customers. Use a third-party routing appliance in the SaaS application VPC to terminate onpremises Site-to-Site VPN connections.

C. Peer each AWS customer's VPCs to the VPC that hosts the SaaS application. Create AWS Site-to- Site VPN connections on the SaaS VPC virtual private gateway.

D. Use Site-to-Site VPN tunnels to connect each AWS customer's VPCs to the VPC that hosts the SaaS application. Use AWS Site-to-Site VPN to connect the on-premises customers.

Answer: B

Explanation:

Question: 164

A company's existing AWS environment contains public application servers that run on Amazon EC2 instances. The application servers run in a VPC subnet. Each server is associated with an Elastic IP address.

The company has a new requirement for firewall inspection of all traffic from the internet before the traffic reaches any EC2 instances. A security engineer has deployed and configured a Gateway Load Balancer (GLB) in a standalone VPC with a fleet of third-party firewalls.

How should a network engineer update the environment to ensure that the traffic travels across the fleet of firewalls?

A. Deploy a transit gateway. Attach a GLB endpoint to the transit gateway. Attach the application VPC to the transit gateway. Update the application subnet route table's default route destination to be the GLB endpoint. Ensure that the EC2 instances' security group allows traffic from the GLB endpoint.

B. Update the application subnet route table to have a default route to the GLB. On the standalone VPC that contains the firewall fleet, add a route in the route table for the application VPC's CIDR block with the GLB endpoint as the destination. Update the EC2 instances' security group to allow traffic from the GLB.

C. Provision a GLB endpoint in the application VPC in a new subnet. Create a gateway route table with a route that specifies the application subnet CIDR block as the destination and the GLB endpoint as the target. Associate the gateway route table with the internet gateway in the application VPC. Update the application subnet route table's default route destination to be the GLB endpoint.

D. Instruct the security engineer to move the GLB into the application VPC. Create a gateway route table. Associate the gateway route table with the application subnet. Add a default route to the gateway route table with the GLB as its destination. Update the route table on the GLB to direct traffic from the internet gateway to the application servers. Ensure that the EC2 instances' security

group allows traffic from the GLB.

Answer: C

Explanation:

Question: 165

A company has an AWS Site-to-Site VPN connection between its office and its VPC. Users report occasional failure of the connection to the application that is hosted inside the VPC. A network engineer discovers in the customer gateway logs that the Internet Key Exchange (IKE) session ends when the connection to the application fails.

What should the network engineer do to bring up the IKE session if the IKE session goes down?

A. Set the dead peer detection (DPD) timeout action to Clear. Initiate traffic from the VPC to on premises.

B. Set the dead peer detection (DPD) timeout action to Restart. Initiate traffic from on premises to the VPC.

C. Set the dead peer detection (DPD) timeout action to None. Initiate traffic from the VPC to on premises.

D. Set the dead peer detection (DPD) timeout action to Cancel. Initiate traffic from on premises to the VPC.

Answer: B

Explanation:

Question: 166

A network engineer is designing a hybrid networking environment that will connect a company's corporate network to the company's AWS environment. The AWS environment consists of 30 VPCs in 3 AWS Regions.

The network engineer needs to implement a solution to centrally filter traffic by using a firewall that the company's security team has approved. The solution must give all the VPCs the ability to connect to each other. Connectivity between AWS and the corporate network must meet a minimum bandwidth requirement of 2 Gbps.

Which solution will meet these requirements?

A. Deploy an IPsec VPN connection between the corporate network and a new transit gateway. Connect all VPCs to the transit gateway. Associate the approved firewall with the transit gateway.

B. Deploy a single 10 Gbps AWS Direct Connect connection between the corporate network and virtual private gateway of each VPC. Connect the virtual private gateways to a Direct Connect gateway. Build an IPsec tunnel to a new transit VPC. Deploy the approved firewall to the transit VPC.

C. Deploy two 1 Gbps AWS Direct Connect connections in different Direct Connect locations to connect to the corporate network.

Build a transit VIF on each connection to a Direct Connect gateway. Associate the Direct Connect gateway with a new transit gateway for each Region. Configure the VIFs to use equal-cost multipath (ECMP) routing. Connect all the VPCs in the three Regions to the transit gateway. Configure the transit gateway route table to route traffic to an inspection VPC. Deploy the approved firewall to the inspection VPC.

D. Deploy four 1 Gbps AWS Direct Connect connections in different Direct Connect locations to connect to the corporate network. Build a transit VIF on each connection to a Direct Connect gateway. Associate the Direct Connect gateway with a new transit gateway for each Region. Connect the transit gateways by using a transit gateway peering attachment. Configure the VIFs to use equalcost multipath (ECMP) routing. Configure transit gateway route tables to route traffic to an inspection VPC. Deploy the approved firewall to the inspection VPC.

Answer: D

Explanation:

Question: 167

A company uses an AWS Direct Connect private VIF with a link aggregation group (LAG) that consists of two 10 Gbps connections. The company's security team has implemented a new requirement for external network connections to provide layer 2 encryption. The company's network team plans to use MACsec support for Direct Connect to meet the new requirement.

Which combination of steps should the network team take to implement this functionality? (Choose three.)

A. Create a new Direct Connect LAG with new circuits and ports that support MACsec.

B. Associate the MACsec Connectivity Association Key (CAK) and the Connection Key Name (CKN) with the new LAG.

C. Associate the Internet Key Exchange (IKE) with the existing LAG.

D. Configure the MACsec encryption mode on the existing LAG.

E. Configure the MACsec encryption mode on the new LAG.

F. Configure the MACsec encryption mode on each Direct Connect connection that makes up the existing LAG.

Answer: A, B, E

Explanation:

Question: 168

A company recently implemented a security policy that prohibits developers from launching VPC network infrastructure. The policy states that any time a NAT gateway is launched in a VPC, the company's network security team must immediately receive an alert to terminate the NAT gateway. The network security team needs to implement a solution that can be deployed across AWS accounts

with the least possible administrative overhead. The solution also must provide the network security team with a simple way to view compliance history.

Which solution will meet these requirements?

A. Develop a script that programmatically checks for NAT gateways in an AWS account, sends an email alert, and terminates the NAT gateway if a NAT gateway is detected. Deploy the script on an Amazon EC2 instance in each account. Use a cron job to run the script every 5 minutes. Log the results of the checks to an Amazon RDS for MySQL database.

B. Create an AWS Lambda function that programmatically checks for NAT gateways in an AWS account, sends an email alert, and terminates the NAT gateway if a NAT gateway is detected. Deploy the Lambda function to each account by using AWS Serverless Application Model (AWS SAM) templates. Store the results of the checks on an Amazon OpenSearch Service cluster in each account.

C. Enable Amazon GuardDuty. Create an Amazon EventBridge rule for the Behavior:EC2/NATGatewayCreation GuardDuty finding type. Configure the rule to invoke an AWS Step Functions state machine to send an email alert and terminate a NAT gateway if a NAT gateway is detected. Store the runtime log as a text file in an Amazon S3 bucket.

D. Create a custom AWS Config rule that checks for NAT gateways in an AWS account. Configure the AWS Config rule to perform an AWS Systems Manager Automation remediation action to send an email alert and terminate the NAT gateway if a NAT gateway is detected. Deploy the AWS Config rule and the Systems Manager runbooks to each account by using AWS CloudFormation StackSets

---

Answer: D

Explanation:

Question: 169

A company is running an online game on AWS. The game is played globally and is gaining popularity. Users are reporting problems with the game's responsiveness. Replay rates are dropping, and the company is losing subscribers. Game servers are located in the us-west-2 Region and use an Elastic Load Balancer to distribute client traffic.

The company has decided to deploy game servers to 11 additional AWS Regions to reduce the roundtrip times of network traffic to game clients. A network engineer must design a DNS solution that uses Amazon Route 53 to ensure that user traffic is delivered to game servers with an optimal response time.

What should the network engineer do to meet these requirements?

A. Create Route 53 records for the Elastic Load Balancers in each Region. Specify a weighted routing policy. Calculate the weight by using the number of clients in each Region.

B. Create Route 53 records for the Elastic Load Balancers in each Region. Specify a latency routing policy. Set the Region to the Region where the Elastic Load Balancer is deployed.\

C. Create Route 53 records for the Elastic Load Balancers in each Region. Specify a multivalue answer routing policy. Test latency from the game client, and connect to the server with the best response.

D. Create Route 53 records for the Elastic Load Balancers in each Region. Specify a geolocation routing policy. Set the location to the Region where the Elastic Load Balancer is deployed.

Answer: B

Explanation:

Question: 170

A network engineer needs to build an encrypted connection between an on-premises data center and a VPC. The network engineer attaches the VPC to a virtual private gateway and sets up an AWS Site-to-Site VPN connection. The VPN tunnel is UP after configuration and is working. However, during rekey for phase 2 of the VPN negotiation, the customer gateway device is receiving different parameters than the parameters that the device is configured to support.

The network engineer checks the IPsec configuration of the VPN tunnel. The network engineer notices that the customer gateway device is configured with the most secure encryption algorithms that the AWS Site-to-Site VPN configuration file provides.

What should the network engineer do to troubleshoot and correct the issue?

A. Check the native virtual private gateway logs. Restrict the VPN tunnel options to the specific VPN parameters that the virtual private gateway requires.

B. Check the native customer gateway logs. Restrict the VPN tunnel options to the specific VPN parameters that the customer gateway requires.

C. Check Amazon CloudWatch logs of the virtual private gateway. Restrict the VPN tunnel options to the specific VPN parameters that the virtual private gateway requires.

D. Check Amazon CloudWatch logs of the customer gateway. Restrict the VPN tunnel options to the specific VPN parameters that the customer gateway requires.

Answer: B

Explanation:

Question: 171

A company is growing rapidly. Data transfers between the company's on-premises systems and Amazon EC2 instances that run in VPCs are limited by the throughput of a single AWS Site-to-Site VPN connection between the company's on-premises data center firewall and an AWS Transit Gateway.

A network engineer must resolve the throttling by designing a solution that is highly available and secure. The solution also must scale the VPN throughput from on premises to the VPC resources to support the increase in traffic.

Which solution will meet these requirements?

A. Configure multiple dynamic BGP-based Site-to-Site VPN connections to the transit gateway Configure equal-cost multi-path routing

B. Configure multiple static routing-based Site-to-Site VPN connections to the transit gateway. Configure equal-cost multi-path routing

C. Configure a new Site-to-Site VPN connection to the transit gateway Enable acceleration for the Site-to-Site VPN connection.

D. Configure a software appliance-based VPN connection over the internet from the on-premises firewall to an EC2 instance that has a large instance size and networking capabilities.

Answer: A

Explanation:

Question: 172

A network engineer is designing the DNS architecture for a new AWS environment. The environment must be able to resolve DNS names of endpoints on premises, and the on-premises systems must be able to resolve the names of AWS endpoints. The DNS architecture must give individual accounts the ability to manage subdomains.

The network engineer needs to create a single set of rules that will work across multiple accounts to control this behavior. In addition, the network engineer must use AWS native services whenever possible.

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

A. Create an Amazon Route 53 private hosted zone for the overall cloud domain. Plan to create subdomains that align to other AWS accounts that are associated with the central Route 53 private hosted zone.

B. Create AWS Directory Service for Microsoft Active Directory server endpoints in the central AWS account that hosts the private hosted zone for the overall cloud domain. Create a conditional forwarding rule in Microsoft Active Directory DNS to forward traffic to a DNS resolver endpoint on premises. Create another rule to forward traffic between subdomains to the VPC resolver.

C. Create Amazon Route 53 Resolver inbound and outbound endpoints in the central AWS account that hosts the private hosted zone for the overall cloud domain. Create a forwarding rule to forward traffic to a DNS resolver endpoint on premises. Create another rule to forward traffic between subdomains to the Resolver inbound endpoint.

D. Ensure that networking exists between the other accounts and the central account so that traffic can reach the AWS Directory Service for Microsoft Active Directory DNS endpoints.

E. Ensure that networking exists between the other accounts and the central account so that traffic can reach the Amazon Route 53 Resolver endpoints.

E. Share the Amazon Route 53 Resolver rules between accounts by using AWS Resource Access Manager (AWS RAM). Ensure that networking exists between the other accounts and the central account so that traffic can reach the Route 53 Resolver endpoints.

Answer: A, C, D

Explanation:

Question: 173

A company wants to migrate its DNS registrar and DNS hosting to Amazon Route 53. The company website receives tens of thousands of visits each day, and the company's current DNS provider cannot keep up. The company wants to migrate as quickly as possible but cannot tolerate any downtime.

Which solution will meet these requirements?

A. Transfer the domain name to Route 53. Create a Route 53 private hosted zone, and copy all the existing DNS records. Update the name servers on the domain to use the name servers that are specified n the newly created private hosted zone.

B. Copy all DNS records from the existing DNS servers to a Route 53 private hosted zone. Update the name servers with the existing registrar to use the private hosted zone name servers. Transfer the domain name to Route 53. Ensure that all the changes have propagated

C. Transfer the domain name to Route 53. Create a Route 53 public hosted zone, and copy all the existing DNS records. Set the TTL value on each record to 1 second. Update the name servers on the domain to use the name servers that are specified in the newly created public hosted zone.

D. Copy all DNS records from the existing DNS servers to a Route 53 public hosted zone. Update the name servers with the existing registrar to use the Route 53 name servers for the hosted zone. When the changes have propagated, perform a domain name transfer to Route 53.

Answer: D

Explanation:

Question: 174

A network engineer needs to provide dual-stack connectivity between a company's office location and an AWS account. The company's on-premises router supports dual-stack connectivity, and the VPC has been configured with dual-stack support. The company has set up two AWS Direct Connect connections to the office location. This connectivity must be highly available and must be reliable for latency-sensitive traffic.

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

A. Configure a single private VIF on each Direct Connect connection. Add both IPv4 and IPv6 peering to each private VIF. Configure the on- premises equipment with the AWS provided BGP neighbors to advertise IPv4 routes on the IPv4 peering and IPv6 routes on the IPv6 peering. Enable Bidirectional Forwarding Detection (BFD) on all peering sessions.

B. Configure two private VIFs on each Direct Connect connection: one private VIF with the IPv4 address family and one private VIF with the IPv6 address family. Configure the on-premises equipment with the AWS provided BGP neighbors to advertise IPv4 routes on the IPv4 peering and IPv6 routes on the IPv6 peering. Enable Bidirectional Forwarding Detection (BFD) on all peering sessions.

C. Configure a single private VIF and IPv4 peering on each Direct Connect connection. Configure the on-premises equipment with this peering to advertise the IPv6 routes in the same BGP neighbor configuration. Enable Bidirectional Forwarding Detection (BFD) on all peering sessions.

D. Configure two private VIFs on each Direct Connect connection: one private VIF with the IPv4 address family and one private VIF with the IPv6 address family. Configure the on-premises equipment with the AWS provided BGP neighbors to advertise all IPv4 routes and IPv6 routes on all peering sessions. Keep the Bidirectional Forwarding Detection (BFD) configuration unchanged.

E. Configure two private VIFs on each Direct Connect connection: one private VIF with the IPv4 address family and one private VIF with the IPv6 address family. Configure the on-premises equipment with the AWS provided BGP neighbors to advertise IPv4 routes on the IPv4 peering and IPv6 routes on the IPv6 peering. Reduce the BGP hello timer to 5 seconds on both the on-premises equipment and the Direct Connect configuration.

Answer: A, B

Explanation:

Question: 175

A company recently started using AWS Client VPN to give its remote users the ability to access resources in multiple peered VPCs and resources in the company's on-premises data center. The Client VPN endpoint route table has a single entry of 0.0.0.0/0. The Client VPN endpoint is using a new security group that has no inbound rules and a single outbound rule that allows all traffic to 0.0.0.0/0.

Multiple users report that web search results are showing remote incorrect geographic location information for the users.

Which combination of steps should a network engineer take to resolve this issue with the LEAST amount of service interruption? (Choose three.)

A. Switch users to AWS Site-to-Site VPNs.

B. Enable the split-tunnel option on the Client VPN endpoint.

C. Add routes for the peered VPCs and for the on-premises data center to the Client VPN route table.

D. Remove the 0.0.0.0/0 outbound rule from the security group that the Client VPN endpoint uses.

E. Delete and recreate the Client VPN endpoint in a different VPC.

F. Remove the 0.0.0.0/0 entry from the Client VPN endpoint route table.

Answer: B, C, F

Explanation:

Question: 176

A company has set up hybrid connectivity between its VPCs and its on-premises data center. The company has the on-premises.example.com subdomain configured at its DNS server in the onpremises data center. The company is using the aws.example.com subdomain for workloads that run on AWS across different VPCs and accounts. Resources in both environments can access each other by using IP addresses. The company wants workloads in the VPCs to be able to access resources on premises by using the on-premises.example.com DNS names.

Which solution will meet these requirements with MINIMUM management of resources?

A. Create an Amazon Route 53 Resolver outbound endpoint. Configure a Resolver rule that conditionally forwards DNS queries for on-premises.example.com to the on-premises DNS server. Associate the rule with the VPCs.

B. Create an Amazon Route 53 Resolver inbound endpoint and a Resolver outbound endpoint. Configure a Resolver rule that conditionally forwards DNS queries for on-premises.example.com to the on-premises DNS server. Associate the rule with the VPCs.

C. Launch an Amazon EC2 instance. Install and configure BIND software to conditionally forward DNS queries for on-premises.example.com to the on-premises DNS server. Configure the EC2 instance's IP address as a custom DNS server in each VPC.

D. Launch an Amazon EC2 instance in each VPC. Install and configure BIND software to conditionally forward DNS queries for on-premises.example.com to the on-premises DNS server. Configure the EC2 instance's IP address as a custom DNS server in each VPC.

Answer: A

Explanation:

Question: 177

A company is in the early stage of AWS Cloud adoption. The company has an application that is running in an on-premises data center in Asia. The company needs to deploy new applications in the us-east-1 Region. The applications in the cloud need connectivity to the on-premises data center.

The company needs to set up a communication channel between AWS and the data center. The solution must improve latency, minimize the possibility of performance impact from transcontinental routing over the public internet, and encrypt data in transit.

Which solution will meet these requirements in the LEAST amount of time?

A. Create an AWS Site-to-Site VPN connection with acceleration turned on. Create a virtual private gateway. Attach the Site-to-Site VPN connection to the virtual private gateway. Attach the virtual private gateway to the VPC where the applications will be deployed.

B. Create an AWS Site-to-Site VPN connection with acceleration turned on. Create a transit gateway. Attach the Site-to-Site VPN connection to the transit gateway. Create a transit gateway attachment to the VPC where the applications will be deployed.

C. Create an AWS Direct Connect connection. Create a virtual private gateway. Create a public VIF and a private VIF that use the virtual private gateway. Create an AWS Site-to-Site VPN connection over the public VIF.

D. Create an AWS Site-to-Site VPN connection with acceleration turned off. Create a transit gateway. Attach the Site-to-Site VPN connection to the transit gateway. Create a transit gateway attachment to the VPC where the applications will be deployed.

Answer: B

Explanation:

Question: 178

A company is moving its record-keeping application to the AWS Cloud. All traffic between the company's on-premises data center and AWS must be encrypted at all times and at every transit device during the migration.

The application will reside across multiple Availability Zones in a single AWS Region. The application will use existing 10 Gbps AWS Direct Connect dedicated connections with a MACsec capable port. A network engineer must ensure that the Direct Connect connection is secured accordingly at every transit device.

The network engineer creates a Connection Key Name and Connectivity Association Key (CKN/CAK) pair for the MACsec secret key.

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

A. Configure the on-premises router with the MACsec secret key.

B. Update the connection's MACsec encryption mode to must_encrypt. Then associate the CKN/CAK pair with the connection.

C. Update the connection's MACsec encryption mode to should encrypt. Then associate the CKN/CAK pair with the connection.

D. Associate the CKN/CAK pair with the connection. Then update the connection's MACsec encryption mode to must_encrypt.

E. Associate the CKN/CAK pair with the connection. Then update the connection's MACsec encryption mode to should_encrypt.

Answer: A, E

Explanation:

Question: 179

A network engineer is designing hybrid connectivity with AWS Direct Connect and AWS Transit Gateway. A transit gateway is attached to a Direct Connect gateway and 19 VPCs across different AWS accounts. Two new VPCs are being attached to the transit gateway. The IP address administrator has assigned 10.0.32.0/21 to the first VPC and 10.0.40.0/21 to the second VPC. The prefix list has one CIDR block remaining before the prefix list reaches the quota for the maximum number of entries.

What should the network engineer do to advertise the routes from AWS to on premises to meet these requirements?

A. Add 10.0.32.0/21 and 10.0.40.0/21 to both AWS managed prefix lists.

B. Add 10.0.32.0/21 and 10.0.40.0/21 to the allowed prefix list.

C. Add 10.0.32.0/20 to both AWS managed prefix lists.

D. Add 10.0.32.0/20 to the allowed prefix list.

Answer: D

Explanation:

Question: 180

A company has a single VPC in the us-east-1 Region. The company is planning to set up a new VPC in the us-east-2 Region. The existing VPC has an AWS Site-to-Site VPN connection to the company's onpremises environment and uses a virtual private gateway.

A network engineer needs to implement a solution to establish connectivity between the existing VPC and the new VPC. The solution also must implement support for IPv6 for the new VPC. The company has new on-premises resources that need to connect to VPC resources by using IPv6 addresses.

Which solution will meet these requirements?

A. Create a new virtual private gateway in us-east-1. Attach the new virtual private gateway to the new VPC. Create two new Site-to-Site VPN connections to the new virtual private gateway with IPv4 and IPv6 support. Configure routing between the VPCs by using VPC peering.

B. Create a transit gateway in us-east-1 and in us-east-2. Attach the existing VPC and the new VPC to each transit gateway. Create a new Site-to-Site VPN connection to each transit gateway with IPv4 and IPv6 support. Configure transit gateway peering. Configure routing between the VPCs and the on premises environment.

C. Create a new virtual private gateway in us-east-2. Attach the new virtual private gateway to the new VPC. Create two new Site-to-Site VPN connections to the new virtual private gateway with IPv4 and IPv6 support. Configure routing between the VPCs by using VPC peering.

D. Create a transit gateway in us-east-1. Attach the existing VPC and the new VPC to the transit gateway. Create two new Site-to-Site VPN connections to the transit gateway with IPv4 and IPv6 support. Configure transit gateway peering. Configure routing between the VPCs and the on premises environment.

Answer: B

Explanation:

Question: 181

A global film production company uses the AWS Cloud to encode and store its video content before distribution. The company's three global offices are connected to the us-east-1 Region through AWS Site-to-Site VPN links that terminate on a transit gateway with BGP routing activated.

The company recently started to produce content at a higher resolution to support 8K streaming. The size of the content files has increased to three times the size of the content files from the previous format. Uploads of files to Amazon EC2 instances are taking 10 times longer than they did with the previous format.

Which actions should a network engineer recommend to reduce the upload times? (Choose two.)

A. Create a second VPN tunnel from each office location to the transit gateway. Activate equal-cost multi-path (ECMP) routing.

B. Modify the transit gateway to activate Jumbo MTU on the VPN tunnels to each office location.

C. Replace the existing VPN tunnels with new tunnels that have acceleration activated.

D. Upgrade each EC2 instance to a modern instance type. Activate Jumbo MTU in the operating system.

E. Replace the existing VPN tunnels with new tunnels that have IGMP activated.

Answer: A, C

Explanation:

Question: 182

An application team for a startup company is deploying a new multi-tier application into the AWS Cloud. The application will be hosted on a fleet of Amazon EC2 instances that run in an Auto Scaling group behind a publicly accessible Network Load Balancer (NLB). The application requires the clients to work with UDP traffic and TCP traffic.

In the near term, the application will serve only users within the same geographic location. The application team plans to extend the application to a global audience and will move the deployment to multiple AWS Regions around the world to bring the application closer to the end users. The application team wants to use the new Regions to deploy new versions of the application and wants to be able to control the amount of traffic that each Region receives during these rollouts. In addition, the application team must minimize first-byte latency and jitter (randomized delay) for the end users.

How should the application team design the network architecture for the application to meet these requirements?

A. Create an Amazon CloudFront distribution to align to each Regional deployment. Set the NLB for each Region as the origin for each CloudFront distribution. Use an Amazon Route 53 weighted routing policy to control traffic to the newer Regional deployments.

B. Create an AWS Global Accelerator accelerator and listeners for the required ports. Configure endpoint groups for each Region. Configure a traffic dial for the endpoint groups to control traffic to the newer Regional deployments. Register the NLBs with the endpoint groups.

C. Use Amazon S3 Transfer Acceleration for the application in each Region. Adjust the amount of traffic that each Region receives from the Transfer Acceleration endpoints to the Regional NLBs.

D. Create an Amazon CloudFront distribution that includes an origin group. Set the NLB for each Region as the origins for the origin group. Use an Amazon Route 53 latency routing policy to control traffic to the new Regional deployments.

Answer: B

Explanation:

Question: 183

A company deploys a software solution on Amazon EC2 instances that are in a cluster placement group. The solution's UI is a single HTML page. The HTML file size is 1,024 bytes. The software processes files that exceed 1,024 MB in size. The software shares files over the network to clients upon request. The files are shared with the Don't Fragment flag set. Elastic network interfaces of the EC2 instances are set up with jumbo frames.

The UI is always accessible from all allowed source IP addresses, regardless of whether the source IP addresses are within a VPC, on the internet, or on premises. However, clients sometimes do not receive files that they request because the files fail to travel successfully from the software to the clients.

Which options provide a possible root cause of these failures? (Choose two.)

A. The source IP addresses are from on-premises hosts that are routed over AWS Direct Connect.

B. The source IP addresses are from on-premises hosts that are routed over AWS Site-to-Site VPN.

C. The source IP addresses are from hosts that connect over the public internet.

D. The security group of the EC2 instances does not allow ICMP traffic.

E. The operating system of the EC2 instances does not support jumbo frames.

Answer: B, C

Explanation:

Question: 184

A company has users who work from home. The company wants to move these users to Amazon WorkSpaces for additional security visibility.

The company has deployed WorkSpaces in its own AWS account in VPC A. A network engineer decides to provide the security visibility by using two firewall appliances behind a Gateway Load Balancer (GWLB). The network engineer provisions another VPC, VPC B, in a separate account and deploys the two firewall appliances in separate Availability Zones.

What should the network engineer do to configure the network connectivity for this solution?

A. Create a GWLB in VPC A with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the WorkSpaces account to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the default route to the VPC endpoint.

B. Create a GWLB in VPC B with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the WorkSpaces account to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the default route to the GWLB endpoint.

C. Create a GWLB in VPC B with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the WorkSpaces account to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the WorkSpaces subnet to the VPC endpoint.

D. Create a GWLB in VPC B with the firewall appliance instances as targets. Use the GWLB to create a GWLB endpoint. Add the AWS principal ARN of the account that contains the firewall appliances to the principal allow list of the GWLB endpoint. In the WorkSpaces account, create a VPC endpoint and specify the service name that the AWS Management Console provides for the GWLB endpoint. Modify the route tables of VPC A to point the default route to the VPC endpoint.

Answer: B

Explanation:

Question: 185

A company plans to run a computationally intensive data processing application on AWS. The data is highly sensitive. The VPC must have no direct internet access, and the company has applied strict network security to control access.

Data scientists will transfer data from the company's on-premises data center to the instances by using an AWS Site-to-Site VPN connection. The on-premises data center uses the network range 172.31.0.0/20 and will use the network range 172.31.16.0/20 in the application VPC.

The data scientists report that they can start new instances of the application but that they cannot transfer any data from the on-premises data center. A network engineer enables VPC flow logs and sends a ping to one of the instances to test reachability. The flow logs show the following:

2 123456789010 eni-1235b8cal23456789 172.31.8.29 172.31.18.139 0 0 1 4 336 1622433184 1622433194 ACCEPT OK 2 123456789010 eni-1235b8eal23456789 172.31.18.139 172.31.8.29 0 0 1 3 252 1622433216 1622433232 REJECT OK

The network engineer must recommend a solution that will give the data scientists the ability to transfer data from the on-premises data center.

Which solution will meet these requirements?

A. Modify the security group for the application. Add an inbound rule to allow traffic from the on-premises data center network range to the application.

B. Modify the network ACLs for the VPC subnet. Add an inbound rule to allow traffic from the on-premises data center network range to the VPC subnet range.

C. Modify the network ACLs for the VPC subnet. Add an outbound rule to allow traffic from the VPC subnet range to the on-premises data center network range.

D. Modify the security group for the application. Add an outbound rule to allow traffic from the application to the on-premises data center network range.

Answer: C

Explanation:

Question: 186

A company needs to temporarily scale out capacity for an on-premises application and wants to deploy new servers on Amazon EC2 instances. A network engineer must design the networking solution for the connectivity and for the application on AWS.

The EC2 instances need to share data with the existing servers in the on-premises data center. The servers must not be accessible from the internet. All traffic to the internet must route through the firewall in the on-premises data center. The servers must be able to access a third-party web application.

Which configuration will meet these requirements?

A. Create a VPC that has public subnets and private subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a NAT gateway in a public subnet. Create a route table, and associate the public subnets with the route table. Add a default route to the internet gateway. Create a route table, and associate the private subnets with the route table. Add a default route to the NAT gateway. Add routes for the data center subnets to the virtual private gateway. Deploy the application to the private subnets.

B. Create a VPC that has private subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a route table, and associate the private subnets with the route table. Add a default route to the virtual private gateway. Deploy the application to the private subnets.

C. Create a VPC that has public subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a route table, and associate the public subnets with the route table. Add a default route to the internet gateway. Add routes for the on-premises data center subnets to the virtual private gateway. Deploy the application to the public subnets.

D. Create a VPC that has public subnets and private subnets. Create a customer gateway, a virtual private gateway, and an AWS Site-to-Site VPN connection. Create a route table, and associate the public subnets with the route table. Add a default route to the internet gateway. Create a route table, and associate the private subnets with the route table. Add routes for the on-premises data center subnets to the virtual private gateway. Deploy the application to the private subnets.

Answer: B

Explanation:

Question: 187

A company is deploying a web application into two AWS Regions. The company has one VPC in each Region. Each VPC has three Amazon EC2 instances as web servers behind an Application Load Balancer (ALB). The company already has configured an Amazon Route 53 public hosted zone for example.com. Users will access the application by using the fully qualified domain name (FQDN) of app.example.com.

The company needs a DNS solution that allows global users to access the application. The solution must route the users' requests to the Region that provides the lowest response time. The solution must fail over to the Region that provides the next-lowest response time if the application is unavailable in the initially intended Region.

Which solution will meet these requirements?

A. For each ALB, create an A record that has a geolocation routing policy to route app.example.com to the IP addresses of the ALB. Configure a Route 53 HTTP health check that monitors each ALB by IP address. Associate the health check with the A records.

B. Create an A record that has a geolocation routing policy to route app.example.com to the IP addresses for both ALBs. Configure a Route 53 health check that monitors TCP port 80 for each ALB by IP address. Associate the health check with the A records.

C. Create an A record that has a latency-based routing policy to route app.example.com as an alias to one of the ALBs. Configure a Route 53 health check that monitors TCP port 80 for each ALB by IP address. Associate the health check with the A records.

D. For each ALB, create an A record that has a latency-based routing policy to route app.example.com as an alias to the ALB. Set the value for Evaluate Target Health to Yes for the records.

Answer: D

Explanation:

Question: 188

A consulting company manages AWS accounts for its customers. One of the company's customers needs to add intrusion prevention for its environment without having to re-architect the environment. The customer's environment includes five VPCs in two AWS Regions in the United States. VPC-to-VPC connectivity is achieved through VPC peering. The customer does not plan to increase the number of VPCs within the next 2 years. The solution must accommodate unencrypted traffic.

Which solution will meet these requirements?

A. Configure VPC security groups and network ACLs.

B. Use an AWS Network Firewall centralized deployment model in each VPC.

C. Use an AWS Network Firewall distributed deployment model in each VPC.

D. Deploy AWS Shield in each VPC.

Answer: C

Explanation:

Question: 189

A marketing company is using hybrid infrastructure through AWS Direct Connect links and a software-defined wide area network (SD-WAN) overlay to connect its branch offices. The company connects multiple VPCs to a third-party SD-WAN appliance transit VPC within the same account by using AWS Site-to-Site VPNs.

The company is planning to connect more VPCs to the SD-WAN appliance transit VPC. However, the company faces challenges of scalability, route table limitations, and higher costs with the existing architecture. A network engineer must design a solution to resolve these issues and remove dependencies.

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

A. Configure a transit gateway to attach the VPCs. Configure a Site-to-Site VPN connection between the transit gateway and the third-party SD-WAN appliance transit VPC. Use the SD-WAN overlay links to connect to the branch offices.

B. Configure a transit gateway to attach the VPCs. Configure a transit gateway Connect attachment for the third-party SD-WAN appliance transit VPC. Use transit gateway Connect native integration of SD-WAN virtual hubs with AWS Transit Gateway.

C. Configure a transit gateway to attach the VPCs. Configure VPC peering between the VPCs and the third-party SD-WAN appliance transit VPC. Use the SD-WAN overlay links to connect to the branch offices.

D. Configure VPC peering between the VPCs and the third-party SD-WAN appliance transit VPC. Use transit gateway Connect native integration of SD-WAN virtual hubs with AWS Transit Gateway.

Answer: B

Explanation:

Question: 190

A company is running a hybrid cloud environment. The company has multiple AWS accounts as part of an organization in AWS Organizations. The company needs a solution to manage a list of IPv4 onpremises hosts that will be allowed to access resources in AWS. The solution must provide version control for the list of IPv4 addresses and must make the list available to the AWS accounts in the organization.

Which solution will meet these requirements?

A. Create a customer-managed prefix list. Add entries for the initial list of on-premises IPv4 hosts. Create a resource share in AWS Resource Access Manager. Add the managed prefix list to the resource share. Share the resource with the organization.

B. Create a customer-managed prefix list. Add entries for the initial list of on-premises IPv4 hosts. Use AWS Firewall Manager to share the managed prefix list with the organization.

C. Create a security group. Add inbound rule entries for the initial list of on-premises IPv4 hosts. Create a resource share in AWS Resource Access Manager. Add the security group to the resource share. Share the resource with the organization.

D. Create an Amazon DynamoDB table. Add entries for the initial list of on-premises IPv4 hosts. Create an AWS Lambda function

that assumes a role in each AWS account in the organization to authorize inbound rules on security groups based on entries from the DynamoDB table.

Answer: A

Explanation:

Question: 191

A company's application is deployed on Amazon EC2 instances in a single VPC in an AWS Region. The EC2 instances are running in two Availability Zones. The company decides to use a fleet of traffic inspection instances from AWS Marketplace to inspect traffic between the VPC and the internet. The company is performing tests before the company deploys the architecture into production.

The fleet is located in a shared inspection VPC behind a Gateway Load Balancer (GWLB). To minimize the cost of the solution, the company deployed only one inspection instance in each Availability Zone that the application uses.

During tests, a network engineer notices that traffic inspection works as expected when the network is stable. However, during maintenance of the inspection instances, the internet sessions time out for some application instances. The application instances are not able to establish new sessions.

Which combination of steps will remediate these issues? (Choose two.)

A. Deploy one inspection instance in the Availability Zones that do not have inspection instances deployed.

B. Deploy one additional inspection instance in each Availability Zone where the inspection instances are deployed.

C. Enable the cross-zone load balancing attribute for the GWLB.

D. Deploy inspection instances in an Auto Scaling group. Define a scaling policy that is based on CPU load.

E. Attach the GWLB to all Availability Zones in the Region.

Answer: B, C

Explanation:

Question: 192

A company deploys an internal website behind an Application Load Balancer (ALB) in a VPC. The VPC has a CIDR block of 172.31.0.0/16. The company creates a private hosted zone for the domain example.com for the website in Amazon Route 53. The company establishes an AWS Site-to-Site VPN connection between its office network and the VPC.

A network engineer needs to set up a DNS solution so that employees can visit the internal webpage by accessing a private domain URL (https://example.com) from the office network.

Which combination of steps will meet this requirement? (Choose two.)

A. Create an alias record that points to the ALB in the Route 53 private hosted zone.

B. Create a CNAME record that points to the ALB internal domain in the Route 53 private hosted zone.

C. Create a Route 53 Resolver inbound endpoint. On the office DNS server, configure a conditional forwarder to forward the DNS queries to the Route 53 Resolver inbound endpoint.

D. Create a Route 53 Resolver outbound endpoint. On the office DNS server, configure a conditional forwarder to forward the DNS queries to the Route 53 Resolver outbound endpoint.

E. On the office DNS server, configure a conditional forwarder for the private domain to the VPC DNS at 172.31.0.2.

Answer: A, C

Explanation:

Question: 193

A company is migrating applications from a data center to AWS. Many of the applications will need to exchange data with the company's on-premises mainframe.

The company needs to achieve 4 Gbps transfer speeds to meet peak traffic demands. A network engineer must design a highly available solution that maximizes resiliency. The solution must be able to withstand the loss of circuits or routers.

Which solution will meet these requirements?

A. Order four 10 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate one connection from each Direct Connect location to a router at the company location. Terminate the other connection from each Direct Connect location to a different router at the company location.

B. Order two 10 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate the connection from each Direct Connect location to a different router at the company location.

C. Order four 1 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate one connection from each Direct Connect location to a router at the company location. Terminate the other connection from each Direct Connect location to a different router at the company location.

D. Order two 1 Gbps AWS Direct Connect connections that are evenly spread over two locations. Terminate the connection from each Direct Connect location to a different router at the company location.

Answer: A

Explanation:

Question: 194

A company has 10 web server Amazon EC2 instances that run in an Auto Scaling group in a production VPC. The company has 10 other web servers that run in an on-premises data center. The company has a 10 Gbps AWS Direct Connect connection between the on-premises data center and the production VPC.

The company needs to implement a load balancing solution that receives HTTPS traffic from thousands of external users. The solution must distribute the traffic across the web servers on AWS and the web servers in the on-premises data center. Regardless of the location of the web servers, HTTPS requests must go to the same web server throughout the entire session.

Which solution will meet these requirements?

A. Create a Network Load Balancer (NLB) in the production VPC. Create a target group. Specify ip as the target type. Register the EC2 instances and the on-premises servers with the target group Enable connection draining on the NLB

B. Create an Application Load Balancer (ALB) in the production VPC. Create a target group Specify ip as the target type. Register the EC2 instances and the on-premises servers with the target group. Enable application-based session affinity (sticky sessions) on the ALB.

C. Create a Network Load Balancer (NLB) in the production VPC. Create a target group. Specify instance as the target type. Register the EC2 instances and the on-premises servers with the target group. Enable session affinity (sticky sessions) on the NLB.

D. Create an Application Load Balancer (ALB) in the production VPC. Create a target group. Specify instance as the target type Register the EC2 instances and the on-premises servers with the target group Enable application-based session affinity (sticky sessions) on the ALB.

Answer: C

Explanation:

Question: 195

A company has an AWS environment that includes multiple VPCs that are connected by a transit gateway. The company has decided to use AWS Site-to-Site VPN to establish connectivity between its on-premises network and its AWS environment.

The company does not have a static public IP address for its on-premises network. A network engineer must implement a solution to initiate the VPN connection on the AWS side of the connection for traffic from the AWS environment to the on-premises network.

Which combination of steps should the network engineer take to establish VPN connectivity between the transit gateway and the on-premises network? (Choose three.)

A. Configure the Site-to-Site VPN tunnel options to use Internet Key Exchange version 1 (IKEv1).

B. Configure the Site-to-Site VPN tunnel options to use Internet Key Exchange version 2 (IKEv2).

C. Use a private certificate authority (CA) from AWS Private Certificate Authority to create a certificate.

D. Use a public certificate authority (CA) from AWS Private Certificate Authority to create a certificate.

E. Create a customer gateway. Specify the current dynamic IP address of the customer gateway device's external interface.

F. Create a customer gateway without specifying the IP address of the customer gateway device.

Answer: B, C, D

Explanation:

Question: 196

A company's AWS environment has two VPCs. VPC A has a CIDR block of 192.168.0.0/16. VPC B has a CIDR block of 10.0.0.0/16. Each VPC is deployed in a separate AWS Region. The company has remote users who work outside the company's offices. These users need to connect to an application that is running in the VPCs.

Traffic to and from the VPCs over the internet must be encrypted. A network engineer must set up connectivity between the remote users and the VPCs.

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

A. Establish an AWS Site-to-Site VPN connection between VPC A and VPC B.

B. Establish a VPC peering connection between VPC A and VPC B.

C. Create an AWS Client VPN endpoint in VPC A and VPC B Add an authorization rule to grant access to VPC A and VPC B.

D. Create an AWS Client VPN endpoint in VPC A Add an authorization rule to grant access to VPC A and VPC B.

E. Add a route to the AWS Client VPN endpoint's route table to direct traffic to VPC B.

F. Add a route to the AWS Client VPN endpoint's route table to direct traffic to VPC A.

Answer: B, D, E

Explanation:

Question: 197

A company uses Amazon Route 53 to register a public domain, example.com, in an AWS account. A central services group manages the account. The company wants to create a subdomain, test.example.com, in another AWS account to offer name services for Amazon EC2 instances that are hosted in the account. The company does not want to migrate the parent domain to the subdomain

account.

A network engineer creates a new Route 53 hosted zone for the subdomain in the second account.

Which combination of steps must the network engineer take to complete the task? (Choose two.)

A. Add records for the hosts of the new subdomain to the new Route 53 hosted zone.

B. Update the DNS service for the parent domain by adding name server (NS) records for the subdomain.

C. Update the DNS service for the subdomain by adding name server (NS) records for the parent domain.

D. Create an alias record from the parent domain that points to the hosted zone for the subdomain in the second account.

E. Add a start of authority (SOA) record in the parent domain for the subdomain.

Answer: A, B

Explanation:

Question: 198

An IoT company collects data from thousands of sensors that are deployed in the Unites States and South Asia. The sensors use a proprietary communication protocol that is built on UDP to send the data to a fleet of Amazon EC2 instances. The instances are in an Auto Scaling group and run behind a Network Load Balancer (NLB). The instances, Auto Scaling group, and NLB are deployed in the us- west-2 Region.

Occasionally, the data from the sensors in South Asia gets lost in transit over the internet and does not reach the EC2 instances.

Which solutions will resolve this issue? (Choose two.)

A. Use AWS Global Accelerator with the existing NLB.

B. Create an Amazon CloudFront distribution. Specify the existing NLB as the origin.

C. Create a second deployment of the EC2 instances and the NLB in the ap-south-1 Region. Use an Amazon Route 53 latency routing policy to resolve to the Region that provides the least latency.

D. Create a second deployment of the EC2 instances and the NLB in the ap-south-1 Region. Use an Amazon Route 53 failover routing policy to resolve to an alternate Region in case packets are dropped.

E. Turn on enhanced networking on the EC2 instances by using the most recent Elastic Network Adapter (ENA) drivers.

Answer: A, C

Explanation:

Question: 199

A company has two AWS Direct Connect links. One Direct Connect link terminates in the us-east-1 Region, and the other Direct Connect link terminates in the af-south-1 Region. The company is using BGP to exchange routes with AWS.

How should a network engineer configure BGP to ensure that af-south-1 is used as a secondary link to AWS?

A.

On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7100

On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7300

On the Direct Connect BGP peer to us-east-1, set the local preference value to 200

On the Direct Connect BGP peer to af-south-1, set the local preference value to 50

B.

On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7300

On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7100

On the Direct Connect BGP peer to us-east-1, set the local preference value to 200

On the Direct Connect BGP peer to af-south-1, set the local preference value to 50

C.

On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7100

On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7300

On the Direct Connect BGP peer to us-east-1, set the local preference value to 50

On the Direct Connect BGP peer to af-south-1, set the local preference value to 200

D.

On the Direct Connect link to us-east-1, configure BGP peering to use community tag 7224:7300

On the Direct Connect link to af-south-1, configure BGP peering to use community tag 7224:7100

On the Direct Connect BGP peer to us-east-1, set the local preference value to 50

On the Direct Connect BGP peer to af-south-1, set the local preference value to 200

Answer: B

Explanation:

Question: 200

A team of infrastructure engineers wants to automate the deployment of Application Load Balancer (ALB) components by using the AWS Cloud Development Kit (AWS CDK). The CDK application must deploy an infrastructure stack that is reusable and consistent across multiple environments, AWS Regions, and AWS accounts.

The lead network architect on the project has already bootstrapped the target accounts. The lead network architect also has deployed core network components such as VPCs and Amazon Route 53 private hosted zones across the multiple environments and Regions. The infrastructure engineers must design the ALB components in the CDK application to use the existing core network components.

Which combination of steps will meet this requirement with the LEAST manual effort between environment deployments? (Choose two.)

A. Design the CDK application to read AWS CloudFormation parameters for the values that vary across environments and Regions. Reference these variables in the CDK stack for resources that require the variables.

B. Design the CDK application to read environment variables that contain account and Region details at runtime. Use these variables as properties of the CDK stack. Use context methods in the CDK stack to retrieve variable values.

C. Create a dedicated account for shared application services in the multi-account environment. Deploy a CDK pipeline to the dedicated account. Create stages in the pipeline that deploy the CDK application across different environments and Regions.

D. Write a script that automates the deployment of the CDK application across multiple environments and Regions. Distribute the script to engineers who are working on the project.

E. Use the CDK toolkit locally to deploy stacks to each environment and Region. Use the --context flag to pass in variables that the CDK application can reference at runtime.

Answer: B, C

Explanation:

Question: 201

A company has critical VPC workloads that connect to an on-premises data center through two redundant active-passive AWS Direct Connect connections. However, a recent outage on one Direct Connect connection revealed that it takes more than a minute for traffic to fail over to the secondary Direct Connect connection. The company wants to reduce the failover time from minutes to seconds.

Which solution will provide the LARGEST reduction in the BGP failover time?

A. Reduce the BGP hold-down timer that is configured on the BGP sessions on the Direct Connect connection VIFs.

B. Configure an Amazon CloudWatch alarm for the Direct Connect connection state to invoke an AWS Lambda function to fail

over the traffic.

C. Configure Bidirectional Forwarding Detection (BFD) on the Direct Connect connections on the AWS side.

D. Configure Bidirectional Forwarding Detection (BFD) on the Direct Connect connections on the onpremises router.

Answer: D

Explanation:

Question: 202

A European car manufacturer wants to migrate its customer-facing services and its analytics platform from two on-premises data centers to the AWS Cloud. The company has a 50-mile (80.4 km) separation between its on-premises data centers and must maintain that separation between its two locations in the cloud. The company also needs failover capabilities between the two locations in the cloud.

The company's infrastructure team creates several accounts to separate workloads and responsibilities. The company provisions resources in the eu-west-3 Region and in the eu-central-1 Region. The company selects an AWS Direct Connect Partner in each Region and requests two resilient 1 Gbps fiber connections from each provider.

The company's network engineer must establish a connection between all VPCs in the accounts and between the on-premises network and the AWS Cloud. The solution must provide access to all services in both Regions in case of network issues.

Which solution will meet these requirements?

A. Create a Direct Connect gateway. Create a private VIF on each of the Direct Connect connections. Attach the private VIFs to the Direct Connect gateway. Use equal-cost multi-path (ECMP) routing to aggregate the four connections across the two Regions. Attach the Direct Connect gateway directly to each VPC's virtual private gateway.

B. Create a Direct Connect gateway. Create a transit gateway. Attach the transit gateway to the Direct Connect gateway. Create a transit VIF on each of the Direct Connect connections. Attach the transit VIFs to the Direct Connect gateway. Use a link aggregation group (LAG) to aggregate the four connections across the two Regions. Attach the transit gateway directly to each VPC.

C. Create a Direct Connect gateway. Create a transit gateway in each Region. Attach the transit gateways to the Direct Connect gateway. Create a transit VIF on each of the Direct Connect connections. Attach the transit VIFs to the Direct Connect gateway. Peer the transit gateways. Attach the transit gateways in each Region to the VPCs in the same Region.

D. Create a Direct Connect gateway. Create a private VIF on each of the Direct Connect connections. Attach the private VIFs to the Direct Connect gateway. Use a link aggregation group (LAG) to aggregate the four connections across the two Regions. Create a transit gateway. Attach the transit gateway to the Direct Connect gateway. Attach the transit gateway directly to each VPC.

Answer: C

Explanation:

Question: 203

A company wants to analyze TCP traffic to the internet. The traffic originates from Amazon EC2 instances in the company's VPC. The EC2 instances initiate connections through a NAT gateway. The required information includes source and destination IP addresses, ports, and the first 8 bytes of payload of TCP segments. The company needs to collect, store, and analyze all the required data points.

Which solution will meet these requirements?

A. Set up the EC2 instances as VPC traffic mirror sources. Deploy software on the traffic mirror target to forward the data to Amazon CloudWatch Logs. Analyze the data by using CloudWatch Logs Insights.

B. Set up the NAT gateway as a VPC traffic mirror source. Deploy software on the traffic mirror target to forward the data to an Amazon OpenSearch Service cluster. Analyze the data by using OpenSearch Dashboards.

C. Turn on VPC Flow Logs on the EC2 instances. Specify the default format and a log destination of Amazon CloudWatch Logs. Analyze the flow log data by using CloudWatch Logs Insights.

D. Turn on VPC Flow Logs on the EC2 instances. Specify a custom format and a log destination of Amazon S3. Analyze the flow log data by using Amazon Athena.

Answer: A

Explanation:

Question: 204

A company has three VPCs in a single AWS Region. Each VPC contains 15 Amazon EC2 instances, and no connectivity exists between the VPCs.

The company is deploying a new application across all three VPCs. The application requires high bandwidth between the nodes. A network engineer must implement connectivity between the VPCs.

Which solution will meet these requirements with the HIGHEST throughput?

A. Configure a transit gateway. Attach each VPC to the transit gateway. Configure static routing in each VPC to route traffic to the transit gateway.

B. Configure VPC peering between the three VPCs. Configure static routing to route traffic between the three VPCs.

C. Configure a transit VPC. Configure a VPN gateway in each VPC. Create an AWS Site-to-Site VPN tunnel from each VPC to the transit VPC. Use BGP routing to route traffic between the VPCs and the transit VPC.

D. Configure AWS Site-to-Site VPN connections between each VPC. Enable route propagation for each Site-to-Site VPN connection to route traffic between the VPCs.

Answer: B

Explanation:

Question: 205

A network engineer needs to deploy an AWS Network Firewall firewall into an existing AWS environment. The environment consists of the following:

A transit gateway with all VPCs attached to it

Several hundred application VPCs

A centralized egress internet VPC with a NAT gateway and an internet gateway

A centralized ingress internet VPC that hosts public Application Load Balancers

On-premises connectivity through an AWS Direct Connect gateway attachment

The application VPCs have workloads deployed across multiple Availability Zones in private subnets with the VPC route table s default route (0.0.0.0/0) pointing to the transit gateway. The Network Firewall firewall needs to inspect east-west (VPC-to-VPC) traffic and north-south (internet-bound and on-premises network) traffic by using Suricata compatible rules.

The network engineer must deploy the firewall by using a solution that requires the least possible architectural changes to the existing production environment.

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

A. Deploy Network Firewall in all Availability Zones in each application VPC.

B. Deploy Network Firewall in all Availability Zones in a centralized inspection VPC.

C. Update the HOME_NET rule group variable to include all CIDR ranges of the VPCs and on-premises networks.

D. Update the EXTERNAL_NET rule group variable to include all CIDR ranges of the VPCs and onpremises networks.

E. Configure a single transit gateway route table. Associate all application VPCs and the centralized inspection VPC with this route table.

F. Configure two transit gateway route tables. Associate all application VPCs with one transit gateway route table. Associate the centralized inspection VPC with the other transit gateway route table.

Answer: B, C, F

Explanation:

Question: 206

A company is using a shared services VPC with two domain controllers. The domain controllers are deployed in the company's private subnets. The company is deploying a new application into a new VPC in the account. The application will be deployed onto an Amazon EC2 for Windows Server instance in the new VPC. The instance must join the existing Windows domain that is supported by the domain controllers in the shared services VPC.

A transit gateway is attached to both the shared services VPC and the new VPC. The company has updated the route tables for the transit gateway, the shared services VPC, and the new VPC. The security groups for the domain controllers and the instance are updated and allow traffic only on the ports that are necessary for domain operations. The instance is unable to join the domain that is hosted on the domain controllers.

Which combination of actions will help identify the cause of this issue with the LEAST operational overhead? (Choose two.)

A. Use AWS Network Manager to perform a route analysis for the transit gateway network. Specify the existing EC2 instance as the source. Specify the first domain controller as the destination. Repeat the route analysis for the second domain controller.\

B. Use port mirroring with the existing EC2 instance as the source and another EC2 instance as the target to obtain packet captures of the connection attempts.

C. Review the VPC flow logs on the shared services VPC and the new VPC.

D. Issue a ping command from one of the domain controllers to the existing EC2 instance.

E. Ensure that route propagation is turned off on the shared services VPC.

Answer: A, C

Explanation:

Question: 207

A company has an order processing system that needs to keep credit card numbers encrypted. The company's customer-facing application runs as an Amazon Elastic Container Service (Amazon ECS) service behind an Application Load Balancer (ALB) in the us-west-2 Region. An Amazon CloudFront distribution is configured with the ALB as the origin. The company uses a third-party trusted certificate authority to provision its certificates.

The company is using HTTPS for encryption in transit. The company needs additional field-level encryption to keep sensitive data encrypted during processing so that only certain application components can decrypt the sensitive data.

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

A. Import the third-party certificate for the ALB. Associate the certificate with the ALB. Upload the certificate for the CloudFront distribution into AWS Certificate Manager (ACM) in us-west-2.

B. Import the third-party certificate for the ALB into AWS Certificate Manager (ACM) in us-west-2. Associate the certificate with the ALB. Upload the certificate for the CloudFront distribution into ACM in the us-east-1 Region.

C. Upload the private key that handles the encryption of the sensitive data to the CloudFront distribution. Create a field-level encryption profile and specify the fields that contain sensitive information. Create a field-level encryption configuration, and choose the newly created profile. Link the configuration to the appropriate cache behavior that is associated with sensitive POST requests.

D. Upload the public key that handles the encryption of the sensitive data to the CloudFront distribution. Create a field-level encryption configuration, and specify the fields that contain sensitive information. Create a field-level encryption profile, and choose the newly created configuration. Link the profile to the appropriate cache behavior that is associated with sensitive GET requests.

E. Upload the public key that handles the encryption of the sensitive data to the CloudFront distribution. Create a field-level encryption profile and specify the fields that contain sensitive information. Create a field-level encryption configuration, and choose the newly created profile. Link the configuration to the appropriate cache behavior that is associated with sensitive POST requests.

Answer: B, E

Explanation:

Question: 208

A company hosts a web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The company uses an Amazon CloudFront distribution with the ALB as an origin.

The application recently experienced an attack. In response, the company associated an AWS WAF web ACL with the CloudFront distribution. The company needs to use Amazon Athena to analyze application attacks that AWS WAF detects.

Which solution will meet this requirement?

A. Configure the ALB and the EC2 instance subnets to produce VPC flow logs. Configure the VPC flow logs to deliver logs to an Amazon S3 bucket for log analysis.

B. Create a trail in AWS CloudTrail to capture data events. Configure the trail to deliver logs to an Amazon S3 bucket for log analysis.

C. Configure the AWS WAF web ACL to deliver logs to an Amazon Kinesis Data Firehose delivery stream. Configure the stream to deliver the data to an Amazon S3 bucket for log analysis.

D. Turn on access logging for the ALB. Configure the access logs to deliver the logs to an Amazon S3 bucket for log analysis.

Answer: D

Explanation:

Question: 209

A real estate company is using Amazon Workspaces to provide corporate managed desktop service to its real estate agents around the world. These Workspaces are deployed in seven VPCs. Each VPC is in a different AWS Region.

According to a new requirement, the company's cloud-hosted security information and events management (SIEM) system needs to analyze DNS queries generated by the Workspaces to identify the target domains that are connected to the Workspaces. The SIEM system supports poll and push methods for data and log collection.

Which solution should a network engineer implement to meet these requirements MOST cost- effectively?

A. Create VPC flow logs in each VPC that is connected to the Workspaces instances. Publish the log data to a central Amazon S3 bucket. Configure the SIEM system to poll the S3 bucket periodically.

B. Configure an Amazon CloudWatch agent to log all DNS requests in Amazon CloudWatch Logs. Configure a subscription filter in CloudWatch Logs. Push the logs to the SIEM system by using Amazon Kinesis Data Firehose.

C. Configure VPC Traffic Mirroring to copy network traffic from each Workspace and to send the traffic to the SIEM system probes for analysis.

D. Configure Amazon Route 53 query logging. Set the destination as an Amazon Kinesis Data Firehose delivery stream that is configured to push data to the SIEM system.

Answer: D

Explanation:

Question: 210

A company uses multiple AWS accounts and VPCs in a single AWS Region. The company must log all network traffic for Amazon EC2 instances and Amazon RDS databases. The company will use the log information to monitor and identify traffic flows in the event of a security incident. The information must be retained for 12 months but will be accessed infrequently after the first 90 days. The company must be able to view metadata that includes the vpc-id, subnet-id: and tcp-flags fields.

Which solution will meet these requirements at the LOWEST cost?

A. Configure VPC flow logs with the default fields Store the logs in Amazon CloudWatch Logs.

B. Configure Traffic Mirroring on all AWS resources to point to a Network Load Balancer that will send the mirrored traffic to monitoring instances.

C. Configure VPC flow logs with additional custom format fields. Store the logs in Amazon S3.

D. Configure VPC flow logs with additional custom format fields. Store the logs in Amazon CloudWatch Logs.

Answer: C

Explanation:

Question: 211

A network engineer is evaluating a network setup for a global retail company. The company has an AWS Direct Connect connection between its on-premises data center and the AWS Cloud. The company has AWS resources in the eu-west-2 Region. These resources consist of multiple VPCs that are attached to a transit gateway.

The company recently provisioned a few AWS resources in the eu-central-1. Region in a single VPC close to its users in this area. The network engineer must connect the resources in eu-central-1 with the on-premises data center and the resources in eu-west-2. The solution must minimize changes to the Direct Connect connection.

What should the network engineer do to meet these requirements?

A. Create a new virtual private gateway. Attach the new virtual private gateway to the VPC in eu-central-1. Use a transit VIF to connect the VPC and the Direct Connect router.

B. Create a new transit gateway in eu-central-1. Create a peering attachment request to the transit gateway in eu-west-2. Add a static route in the transit gateway route table in eu-central-1 to point to the transit gateway peering attachment. Accept the peering request. Add a static route in the transit gateway route table in eu-west-2 to point to the new transit gateway peering attachment.

C. Create a new transit gateway in eu-central-1. Use an AWS Site-to-Site VPN connection to peer both transit gateways. Add a static route in the transit gateway route table in eu-central-1 to point to the transit gateway VPN attachment. Add a static route in the transit gateway route table in eu-west-2 to point to the new transit gateway peering attachment.

D. Create a new virtual private gateway. Attach the new virtual private gateway to the VPC in eu-central-1. Use a public VIF to connect the VPC and the Direct Connect router.

Answer: B

Explanation:

Question: 212

A financial company that is located in the us-east-1 Region needs to establish secure connectivity to AWS. The company has two on-premises data centers, each located within the same Region. The company's network team needs to establish hybrid connectivity to

its AWS environment with reliable and consistent connectivity.

The connection must provide access to the company's private resources inside its AWS environment. The resources are located in the us-east-1 and us-west-2 Regions. The connection must allow resources from the corporate networks to send large amounts of data to Amazon S3 over the same connection. To meet compliance requirements, the connection must be highly available and must provide encryption for all packets that are sent between the on-premises location and any services on AWS.

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

A. Set up a private VIF to send data to Amazon S3. Use an AWS Site-to-Site VPN connection over the private VIF to encrypt data in transit to the VPCs in us-east-1 and us-west-2.

B. Set up an AWS Direct Connect connection to each of the company's data centers.

C. Set up an AWS Direct Connect connection from one of the company's data centers to us-east-1 and us-west-2.

D. Set up a public VIF to send data to Amazon S3. Use an AWS Site-to-Site VPN connection over the public VIF to encrypt data in transit to the VPCs in us-east-1 and us-west-2.

E. Set up a transit VIF for an AWS Direct Connect gateway to send data to Amazon S3. Create a transit gateway. Associate the transit gateway with the Direct Connect gateway to provide secure communications from the company's data centers to the VPCs in us-east-1 and us-west-2.

Answer: B, D

Explanation:

Question: 213

A global company is designing a hybrid architecture to privately access AWS resources in the us-west- 2 Region. The company's existing architecture includes a VPC that uses RFC 1918 IP address space.

The VPC is connected to an on-premises data center over AWS Direct Connect Amazon Route 53 provides name resolution within the VPC. Locally managed DNS servers in the data center provide DNS services to the on-premises hosts.

The company has applications in the data center that need to download objects from an Amazon S3 bucket in us-west-2.

Which solution can the company use to access Amazon S3 without using the public IP address space?

A. Create an S3 interface endpoint in the VPC. Update the on-premises application configuration to use the Regional VPC endpoint DNS hostname that is mapped to the S3 interface endpoint.

B. Create an S3 interface endpoint in the VPC. Configure a Route 53 Resolver inbound endpoint in the VPC. Set up the data center DNS servers to forward DNS queries for the S3 domain from on premises to the inbound endpoint.

C. Create an S3 gateway endpoint in the VPC. Update the on-premises application configuration to use the hostname that is mapped to the S3 gateway endpoint.

D. Create an S3 gateway endpoint in the VPC. Configure a Route 53 Resolver inbound endpoint in the VPC. Set up the data center DNS servers to forward DNS queries for the S3 domain from on premises to the inbound endpoint.

Answer: C

Explanation:

Question: 214

A company is migrating critical applications to AWS. The company has multiple accounts and VPCs that are connected by a transit gateway.

A network engineer must design a solution that performs deep packet inspection for any traffic that leaves a VPC network boundary. All inspected traffic and the actions that are taken on the traffic must be logged in a central log account.

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

A. Create a central network VPC that includes an attachment to the transit gateway. Update the VPC and transit gateway route tables to support the new attachment. Deploy an AWS Gateway Load Balancer that is backed by third-party, next-generation firewall appliances to the central network VPC. Create a policy that contains the rules for deep packet inspection. Attach the policy to the firewall appliances. Create an Amazon S3 bucket in the central log account. Configure the firewall appliances to capture and save the network flow logs to the S3 bucket.

B. Create a central network VPC that includes an attachment to the transit gateway. Update the VPC and transit gateway route tables to support the new attachment. Deploy an AWS Application Load Balancer that is backed by third-party, next-generation firewall appliances to the central network VPC. Create a policy that contains the rules for deep packet inspection. Attach the policy to the firewall appliances. Create a syslog server in the central log account. Configure the firewall appliances to capture and save the network flow logs to the syslog server.

C. Deploy network ACLs and security groups to each VPC. Attach the security groups to active network interfaces. Associate the network ACLs with VPC subnets. Create rules for the network ACLs and security groups to allow only the required traffic flows between subnets and network interfaces. Create an Amazon S3 bucket in the central log account. Configure a VPC flow log that captures and saves all traffic flows to the S3 bucket.

D. Create a central log VPC and an attachment to the transit gateway. Update the VPC and transit gateway route tables to support the new attachment. Deploy an AWS Network Load Balancer (NLB) that is backed by third-party, next-generation intrusion detection system (IDS) security appliances to the central VPC. Activate rules on the security appliances to monitor for intrusion signatures. For each network interface, create a VPC Traffic Mirroring session that sends the traffic to the central VPC's NLB.

Answer: A

Explanation:

Question: 215

A company has an on-premises data center in the United States. The data center is connected to AWS by an AWS Direct Connect connection. The data center has a private VIF that is connected to a Direct Connect gateway.

Recently, the company opened a new data center in Europe and established a new Direct Connect connection between the Europe data center and AWS. A new private VIF connects to the existing Direct Connect gateway.

The company wants to use Direct Connect SiteLink to set up a private network between the data center in the United States and the data center in Europe.

Which solution will meet these requirements in the MOST operationally efficient manner?

A. Create a new public VIF from each data center. Enable SiteLink on the new public VIFs.

B. Create a new transit VIF from each data center. Enable SiteLink on the new transit VIFs.

C. Use the existing VIF from each data center. Enable SiteLink on the existing private VIFs.

D. Create a new AWS Site-to-Site VPN connection between the data centers. Configure the new connection to use SiteLink.

Answer: C

Explanation:

Question: 216

A company has a new AWS Direct Connect connection between its on-premises data center and the AWS Cloud. The company has created a new private VIF on this connection. However, the VIF status is DOWN.

A network engineer verifies that the physical connection status is UP and RUNNING based on information from the AWS Management Console. The network engineer checks the customer Direct Connect router and can see the ARP entry for the VLAN interface created for the private VIF at AWS.

What could be causing the private VIF to have a DOWN status?

A. ICMP is blocked on the customer Direct Connect router.

B. TCP port 179 is blocked on the customer Direct Connect router.

C. The IEEE 802.1Q VLAN identifier is misconfigured on the customer Direct Connect router.

D. The company has configured IEEE 802.1ad instead of 802.1Q on the customer Direct Connect router.

Answer: B

Explanation:

Question: 217

A company recently experienced an IP address exhaustion event in its VPCs. The event affected service capacity. The VPCs hold two or more subnets in different Availability Zones.

A network engineer needs to develop a solution that monitors IP address usage across resources in the VPCs. The company needs to receive notification about possible issues so that the company can act before an incident happens.

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

A. Set up Amazon VPC IP Address Manager (IPAM) with a new top-level pool. In the top-level pool, create a pool for each VPC. In each VPC pool, create a pool for each subnet in that VPC. Turn on the auto-import option for the VPC pools and the subnet pools. Configure an Amazon CloudWatch alarm to send an Amazon Simple Notification Service (Amazon SNS) notification if the availability limit threshold is reached.

B. Set up a log group in Amazon CloudWatch Logs for each subnet. Create an AWS Lambda function that reads each subnet's IP address usage and publishes metrics to the log group. Configure an Amazon CloudWatch alarm to send an Amazon Simple Notification Service (Amazon SNS) notification if the availability limit threshold is reached.

C. Set up a custom Amazon CloudWatch metric for IP address usage for each subnet. Create an AWS Lambda function that reads each subnet's IP address usage and publishes a CloudWatch metric dimension. Schedule the Lambda function to run every 5 minutes. Configure a CloudWatch alarm to send an Amazon Simple Notification Service (Amazon SNS) notification if the availability limit threshold is reached.

D. Set up Amazon VPC IP Address Manager (IPAM) with a new top-level pool. In the top-level pool, create a pool for each VPC. In each VPC pool, create a pool for each subnet in that VPC. Turn on the auto-import option for the VPC pools and the subnet pools. Configure an Amazon EventBridge rule that monitors each pool availability limit threshold and sends an Amazon Simple Notification Service

(Amazon SNS) notification if the limit threshold is reached.

Answer: A

Explanation:

Question: 218

A company is developing a new application that is deployed in multiple VPCs across multiple AWS Regions. The VPCs are connected through AWS Transit Gateway. The VPCs contain private subnets and public subnets.

All outbound internet traffic in the private subnets must be audited and logged. The company's network engineer plans to use AWS Network Firewall and must ensure that all traffic through Network Firewall is completely logged for auditing and alerting.

How should the network engineer configure Network Firewall logging to meet these requirements?

A. Configure Network Firewall logging in Amazon CloudWatch to capture all alerts. Send the logs to a log group in Amazon CloudWatch Logs.

B. Configure Network Firewall logging in Network Firewall to capture all alerts and flow logs.

C. Configure Network Firewall logging by configuring VPC Flow Logs for the firewall endpoint. Send the logs to a log group in Amazon CloudWatch Logs.

D. Configure Network Firewall logging by configuring AWS CloudTrail to capture data events.

Answer: B

Explanation:

Question: 219

A company has set up a NAT gateway in a single Availability Zone (AZ1) in a VPC (VPC1) to access the internet from Amazon EC2 workloads in the VPC. The EC2 workloads are running in private subnets in three Availability Zones (AZ1, AZ2, AZ3). The route table for each subnet is configured to use the NAT gateway to access the internet.

Recently during an outage, internet access stopped working for the EC2 workloads because of the NAT gateway's unavailability. A network engineer must implement a solution to remove the single point of failure from the architecture and provide built-in redundancy.

Which solution will meet these requirements?

A. Set up two NAT gateways. Place each NAT gateway in a different public subnet in separate Availability Zones (AZ2 and AZ3). Configure a route table for private subnets to route traffic to the virtual IP addresses of the two NAT gateways.

B. Set up two NAT gateways. Place each NAT gateway in a different public subnet in separate Availability Zones (AZ2 and AZ3). Configure a route table to point the AZ2 private subnets to the NAT gateway in AZ2. Configure the same route table to point the AZ3 private subnets to the NAT gateway in AZ3.

C. Create a second VPC (VPC2). Set up two NAT gateways. Place each NAT gateway in a different VPC (VPC1 and VPC2) and in the same Availability Zone (AZ2). Configure a route table in VPC1 to point the AZ2 private subnets to one NAT gateway. Configure a route table in VPC2 to point the AZ2 private subnets to the second NAT gateway.

D. Set up two NAT gateways. Place each NAT gateway in a different public subnet in separate Availability Zones (AZ2 and AZ3). Configure a route table to point the AZ2 private subnets to the NAT gateway in AZ2. Configure a second route table to point the AZ3 private subnets to the NAT gateway in AZ3.

Answer: D

Explanation:

Question: 220

A company has agreed to collaborate with a partner for a research project. The company has multiple VPCs in the us-east-1 Region that use CIDR blocks within 10.10.0.0/16. The VPCs are connected by a transit gateway that is named TGW-C in us-east-1. TGW-C has an Autonomous System Number (ASN) configuration value of 64520.

The partner has multiple VPCs in us-east-1 that use CIDR blocks within 172.16.0.0/16. The VPCs are connected by a transit gateway that is named TGW-P in us-east-1. TGW-P has an ASN configuration value of 64530.

A network engineer needs to establish network connectivity between the company's VPCs and the partner's VPCs in us-east-1.

Which solution will meet these requirements with MINIMUM changes to both networks?

A. Create a new VPC in a new account. Deploy a router from AWS Marketplace. Share TGW-C and TGW-P with the new account by using AWS Resource Access Manager (AWS RAM). Associate TGW-C and TGW-P with the new VPC. Configure the router in the new VPC to route between TGW-C and TGW-P.

B. Create an IPsec VPN connection between TGW-C and TGW-P. Configure the routing between the transit gateways to use the IPsec VPN connection.

C. Configure a cross-account transit gateway peering attachment between TGW-C and TGW-P.
Configure the routing between the transit gateways to use the peering attachment.

D. Share TGW-C with the partner account by using AWS Resource Access Manager (AWS RAM).
Associate the partner VPCs with TGW-C. Configure routing in the partner VPCs and TGW-C.

Answer: C

Explanation:

Question: 221

A company has a public application. The application uses an Application Load Balancer (ALB) that has a target group of Amazon EC2 instances.

The company wants to protect the application from security issues in web requests. The traffic to the application must have end-to-end encryption.

Which solution will meet these requirements?

A. Configure a Network Load Balancer (NLB) that has a target group of the existing EC2 instances. Configure TLS connections to terminate on the EC2 instances that use a public certificate. Configure an AWS WAF web ACL. Associate the web ACL with the NLB.

B. Configure TLS connections to terminate at the ALB that uses a public certificate. Configure AWS Certificate Manager (ACM) certificates for the communication between the ALB and the EC2 instances. Configure an AWS WAF web ACL. Associate

the web ACL with the ALB.

C. Configure a Network Load Balancer (NLB) that has a target group of the existing EC2 instances. Configure TLS connections to terminate at the EC2 instances by creating a TLS listener. Configure selfsigned certificates on the EC2 instances for the communication between the NLB and the EC2 instances. Configure an AWS WAF web ACL. Associate the web ACL with the NLB.

D. Configure a third-party certificate on the EC2 instances for the communication between the ALB and the EC2 instances. Import the third-party certificate into AWS Certificate Manager (ACM). Associate the imported certificate with the ALB. Configure TLS connections to terminate at the ALB. Configure an AWS WAF web ACL. Associate the web ACL with the ALB.

Answer: D

Explanation:

Question: 222

A company has an application that hosts personally identifiable information (PII) of users. All connections to the application must be secured by HTTPS with TLS certificates that implement Elliptic Curve Cryptography (ECC).

The application uses stateful connections between the web tier and the end users. Multiple instances host the application. A network engineer must implement a solution that ofloads TLS connections to a load balancer.

Which load-balancing solution will meet these requirements?

A. Provision a Network Load Balancer. Configure a TLS listener by specifying the use of an ECC SSL certificate that is uploaded to AWS identity and Access Management (IAM). Turn on health checks to monitor the web hosts that connect to the end users.

B. Provision an Application Load Balancer. Configure an HTTPS listener by specifying the use of an ECC SSL certificate that is uploaded to AWS Certificate Manager (ACM). Configure a default action to redirect to the URL for the application. Turn on health checks to monitor the web hosts that connect to the end users.

C. Provision a Network Load Balancer. Configure a TLS listener by specifying the use of an ECC SSL certificate that is uploaded to AWS Certificate Manager (ACM). Turn on application-based session affinity (sticky sessions). Turn on health checks to monitor the web hosts that connect to the end users.

D. Provision an Application Load Balancer. Configure an HTTPS listener by specifying the use of an ECC SSL certificate that is uploaded to AWS Identity and Access Management (IAM). Configure a default action to redirect to the URL for the application. Turn on application-based session affinity (sticky sessions).

Answer: D

Explanation:

Question: 223

A company hosts infrastructure services in multiple VPCs across multiple accounts in the us-west-2 Region. The VPC CIDR blocks do not overlap. The company wants to connect the VPCs to its data centers by using AWS Site-to-Site VPN tunnels.

The connections must be encrypted in transit. Additionally, the connection from each data center must route to the closest AWS edge location. The connections must be highly available and must accommodate automatic failover.

Which solution will meet these requirements?

A. Deploy a transit gateway. Share the transit gateway with each of the other accounts by using AWS Resource Access Manager (AWS RAM). Create VPC attachments to the transit gateway from each service account. Add routes to the on-premises subnet in each of the service VPC route tables by using the attachment as the gateway. Create Site-to-Site VPN tunnel attachments with dynamic routing to the transit gateway. Enable the acceleration feature for the Site-to-Site VPN connection.

Configure the VPN tunnels on the on-premises equipment. Configure BGP peering.

B. Deploy VPN gateways to each account. Enable the acceleration feature for VPN gateways on each account. Add routes to the on-premises subnet in each of the service VPC route tables. Use the VPNs as the gateway. Configure the VPN tunnels on the on-premises equipment. Configure BGP peering.

C. Deploy a transit gateway. Share the transit gateway with each of the other accounts by using AWS Resource Access Manager (AWS RAM). Create VPC attachments to the transit gateway from each service account. Add routes to the on-premises subnet in each of the service VPC route tables by using the attachment as the gateway. Create Site-to-Site VPN tunnel attachments with dynamic routing to the transit gateway. Enable the acceleration feature for the Site-to-Site VPN connection. Configure the VPN tunnels on the on-premises equipment. Configure static routing.

D. Deploy VPN gateways to each account. Enable the acceleration feature for VPN gateways on each account. Add routes to the on-premises subnet in each of the service VPC route tables. Use the VPNs as the gateway. Configure the VPN tunnels on the on-premises equipment. Configure static routing.

Answer: A

Explanation:

Question: 224

A company has a transit gateway in AWS Account A. The company uses AWS Resource Access Manager (AWS RAM) to share the transit gateway so that users in other accounts can connect to multiple VPCs in the same AWS Region. AWS Account B contains a VPC (10.0.0.0/16) with subnet 10.0.0.0/24 in the us-west-2a Availability Zone and subnet 10.0.1.0/24 in the us-west-2b Availability Zone. Resources in these subnets can communicate with other VPCs.

A network engineer creates two new subnets: 10.0.2.0/24 in the us-west-2b Availability Zone and 10.0.3.0/24 in the us-west-2c Availability Zone. All the subnets share one route table. The default route 0.0.0.0/0 is pointing to the transit gateway. Resources in subnet 10.0.2.0/24 can communicate with other VPCs, but resources in subnet 10.0.3.0/24 cannot communicate with other VPCs.

What should the network engineer do so that resources in subnet 10.0.3.0/24 can communicate with other VPCs?

A. In Account B, add 10.0.2.0/24 and 10.0.3.0/24 as the destinations to the route table. Use the transit gateway as the target.

B. In Account B, update the transit gateway attachment. Attach the new subnet ID that is associated with us-west-2c to Account B's VPC.

C. In Account A, create a static route for 10.0.3.0/24 in the transit gateway route tables.

D. In Account A, recreate propagation for 10.0.0.0/16 in the transit gateway route tables.

Answer: B

Explanation:

Question: 225

A company has started using AWS Cloud WAN with one edge location in the us-east-1 Region. The company has a production segment and a security segment in AWS Cloud WAN. The company also has a default core network policy.

The company has created a production VPC for the production workload. The company has created an outbound inspection VPC to inspect internet-bound traffic from the production VPC. The company has attached the production VPC to the production segment and has attached the outbound inspection VPC to the security segment. The company has also created an AWS Network Firewall firewall in the outbound inspection VPC to inspect internet-based traffic.

The company has updated a route table for the production VPC to send all internet-bound traffic to the AWS Cloud WAN core network. The company has updated a route table for the outbound inspection VPC to ensure that Network Firewall inspects any outgoing traffic and incoming traffic.

During testing, an Amazon EC2 instance in the production VPC cannot reach the internet. The company checks the Network Firewall rules and confirms that the rules are not blocking the traffic.

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

A. Update the core network policy to configure segment sharing. Share the production segment with the security segment.

B. Update the core network policy to create a static route for the security segment. Specify 0.0.0.0/0 as the destination CIDR block. Specify the outbound inspection VPC as an attachment.

C. Update the core network policy to create a static route for the production segment. Specify 0.0.0.0/0 as the destination CIDR block. Specify the outbound inspection VPC as an attachment.

D. Update the core network policy to create a static route for the production segment. Specify 10.2.0.0/16 as the destination CIDR block. Specify the outbound inspection VPC as an attachment.

E. Create an attachment to attach the outbound inspection VPC to the production segment. Update the core network policy to

turn on isolated attachment for the production segment.

Answer: A, C

Explanation:

Question: 226

A company has an AWS Site-to-Site VPN connection between AWS and its branch office. A network engineer is troubleshooting connectivity issues that the connection is experiencing. The VPN connection terminates at a transit gateway and is statically routed. In the transit gateway route table, there are several static route entries that target specific subnets at the branch office.

The network engineer determines that the root cause of the issues was the expansion of underlying subnet ranges in the branch office during routine maintenance.

Which solution will solve this problem with the LEAST administrative overhead for future expansion efforts?

A. Determine a supernet for the branch office. In the transit gateway route table, add an aggregate route that targets the VPN attachment. Replace the specific subnet routes in the transit gateway route table with the new supernet route.

B. Create an AWS Direct Connect gateway and a transit VIF. Associate the Direct Connect gateway with the transit gateway. Create a propagation for the Direct Connect attachment to the transit gateway route table.

C. Create a dynamically routed VPN connection on the transit gateway. Connect the dynamically routed VPN connection to the branch office. Create a propagation for the VPN attachment to the transit gateway route table. Remove the existing static VPN connection.

D. Create a prefix list that contains the new subnets and the old subnets for the branch office. Remove the specific subnet routes in the transit gateway route table. Create a prefix list reference in the transit gateway route table.

Answer: C

Explanation:

Question: 227

An education agency is preparing for its annual competition between schools. In the competition, students at schools from around the country solve math problems, complete puzzles, and write essays.

The IP addressing plan of all the schools is well-known and is administered centrally. The competition is hosted in the AWS Cloud and is not publicly available. All competition traffic must be encrypted in transit. Only authorized endpoints can access the competition. All the schools have firewall policies that block ICMP traffic.

A network engineer builds a solution in which all the schools access the competition through AWS Site-to-Site VPN connections. The network engineer uses BGP as the routing protocol. The network engineer must implement a solution that notifies schools when they lose connectivity and need to take action on their premises to address the issue.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

A. Monitor the state of the VPN tunnels by using Amazon CloudWatch. Create a CloudWatch alarm that uses Amazon Simple Notification Service (Amazon SNS) to notify people at the affected school if the tunnels are down.

B. Create a scheduled AWS Lambda function that pings each school's on-premises customer gateway device. Configure the Lambda function to send an Amazon Simple Notification Service (Amazon SNS) notification to people at the affected school if the ping fails.

C. Create a scheduled AWS Lambda function that uses the VPC Reachability Analyzer API to verify the connectivity. Configure the Lambda function to send an Amazon Simple Notification Service (Amazon SNS) notification to people at the affected school if failure occurs.

D. Create an Amazon CloudWatch dashboard for each school to show all CloudWatch metrics for each school's Site-to-Site VPN connection. Share each dashboard with the appropriate school.

E. Create a scheduled AWS Lambda function to monitor the existence of each school's routes in the VPC route table where VPN routes are propagated. Configure the Lambda function to send an Amazon Simple Notification Service (Amazon SNS) notification to people at the affected school if failure occurs.

Answer: A, E

Explanation:

Question: 228

A company securely connects resources that are in its VPC to a software as a service (SaaS) solution from a SaaS provider. The SaaS solution is hosted in the AWS Cloud and is powered by AWS PrivateLink. The company uses a PrivateLink endpoint to access the SaaS solution behind the SaaS provider's Network Load Balancer (NLB).

The company recently added a new Availability Zone and new subnets to its VPC. A network engineer is unable to deploy a new interface VPC endpoint for the SaaS solution in the new Availability Zone.

What is the cause of this problem?

A. The CIDR block of the new subnets conflicts with the SaaS provider's CIDR block.

B. The enableDnsHostnames attribute and enableDnsSupport attribute were not configured on the new subnets in the new Availability Zone.

C. The SaaS provider does not offer the solution in the new Availability Zone and has not configured cross-zone load balancing for the NLB.

D. The new subnets are missing a route to the VPC internet gateway.

Answer: C

Explanation:

Question: 229

A company wants to use an AWS Network Firewall firewall to secure its workloads in the cloud through network traffic inspection. The company must record complete metadata information, such as source/destination IP addresses and protocol type. The company must also record all network traffic flows and any DROP or ALERT actions that the firewall takes for traffic that the firewall processes. The Network Firewall endpoints are placed in the correct subnets, and the VPC route tables direct traffic to the Network Firewall endpoints on the path to and from the internet.

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

A. Create a firewall policy to ensure that traffic is processed by stateless or stateful rules according to needs. Select Amazon CloudWatch Logs as the destination for the flow logs.

B. Create a firewall policy to ensure that traffic is processed by stateless or stateful rules according to needs. Configure Network Firewall logging for alert logs and flow logs.

C. Select a destination for logs separately for stateful and stateless engines.

D. Create a firewall policy to ensure that a stateful engine processes all the traffic. Configure Network Firewall logging for alert logs and flow logs. Select a destination for alert logs and flow logs.

E. Create a firewall policy to ensure that a stateful engine processes all the traffic. Configure VPC flow logs for the subnets that the firewall protects. Select a destination for the flow logs.

Answer: C

Explanation:

Question: 230

A network engineer configures a second AWS Direct Connect connection to an existing network. The network engineer runs a test in the AWS Direct Connect Resiliency Toolkit on the connections. The test produces a failure. During the failover event, the network engineer observes a 90-second interruption before traffic shifts to the failover connection.

Which solution will reduce the time for failover?

A. Decrease the BGP hello timer to 5 seconds.

B. Add a VPN connection to the connectivity solution. Implement fast failover.

C. Configure Bidirectional Forwarding Detection (BFD) on the on-premises router.

D. Decrease the BGP hold-down timer to 5 seconds.

Answer: C

Explanation:

Question: 231

A company is building an API-based application on AWS and is using a microservices architecture for the design. The company is using a multi-account AWS environment that includes a separate AWS account for each microservice development team. Each team hosts its microservice in its own VPC that contains Amazon EC2 instances behind a Network Load Balancer (NLB).

A network engineer needs to use Amazon API Gateway in a shared services account to create an HTTP API to expose these microservices to external applications. The network engineer must ensure that access to the microservices can occur only over a private network. Additionally, the company must be able to control which entities from its internal network can connect to the microservices. In the future, the company will create more microservices that the company must be able to integrate with the application.

What is the MOST secure solution that meets these requirements?

A. Create an Application Load Balancer (ALB) in a VPC in the shared services account. Configure the integration to the API Gateway API by using a VPC link. Associate the VPC link with the ALB. Create a VPC endpoint service in each microservice account. Create an AWS PrivateLink endpoint for those services in the shared services account. Add the elastic network interface IP addresses of the VPC endpoint as targets for the target group of the ALB.

B. Create an Application Load Balancer (ALB) in a VPC in the shared services account. Configure the integration to the API Gateway API by using a VPC link. Associate the VPC link with the ALB. Connect all the VPCs to each other by using a central transit gateway. Add the IP addresses of the NLB as IPbased targets in the ALB target group.

C. Configure the integration to the API Gateway API by using HTTP-based integration. Connect all the VPCs to each other by using a central transit gateway. Create a separate HTTP integration to each NLB for each microservice. Add the HTTP endpoint of the NLB as the endpoint URL in the HTTP integration.

D. Configure the integration to the API Gateway API by using VPC link integration. Connect all the VPCs to each other by using a central transit gateway. Create a separate VPC link to each NLB for each microservice. Add the HTTP endpoint of the NLB as the endpoint URL in the VPC link integration.

Answer: A

Explanation:

Question: 232

An international company wants to implement a multi-site hybrid infrastructure. The company wants to deploy its cloud computing resources on AWS in the us-east-1 Region and in the eu-west-2 Region, and in on-premises data centers in the United States (US) and in the United Kingdom (UK). The data centers are connected to each other by a private WAN connection. IP routing information is exchanged dynamically through BGP. The company wants to have two AWS Direct Connect connections, one each in the US and the UK.

The company expects to have 15 VPCs in each Region with CIDR blocks that do not overlap with each other or with CIDR blocks of the on-premises environment. The VPC CIDR blocks are planned so that the prefix aggregation can be performed both on a Regional level and across the entire AWS environment. The company will deploy a transit gateway in each Region to connect the VPCs. A network engineer plans to use a Direct Connect gateway in each Region. A transit VIF will attach the Direct Connect gateway in each Region to the transit gateway in that Region. The transit gateways will be peered with each other.

The network engineer wants to ensure that traffic follows the shortest geographical path from source to destination. Traffic between the on-premises data centers and AWS must travel across a local Direct Connect connection. Traffic between the US data center and eu-west-2 and traffic between the UK data center and us-east-1 must use the private WAN connection to reach the Direct Connect connection to the appropriate Region when the Direct Connect connection is available. The network must be resilient to failures in either the private WAN connection or with the Direct Connect connections. The network also must reroute traffic automatically in the event of any failure.

How should the network engineer configure the transit VIF associations on the Direct Connect gateways to meet these requirements?

A. Advertise only the aggregate route for the company's entire AWS environment.

B. Advertise VPC-specific CIDR prefixes from only the local Region. Additionally, advertise the aggregate route for the company's entire AWS environment.

C. Advertise all the specific VPC CIDR blocks from both Regions.

D. Advertise both Regional aggregate prefixes. Configure custom BGP communities on the routes advertised toward the data center.

Answer: B

Explanation:

Question: 233

Company A recently acquired Company B. Company A has a hybrid AWS and on-premises environment that uses a hosted AWS Direct Connect connection, a Direct Connect gateway, and a transit gateway. Company A has a transit VIF to access the resources in its production environment in the us-east-1 Region.

Company B has applications that run across multiple VPCs in the us-west-2 Region in a single AWS account. A transit gateway connects all Company B's application VPCs. The CIDR blocks for both companies do not overlap.

Company A needs to use the existing Direct Connect connection to access Company B's applications from the on-premises environment.

Which solution will meet these requirements?

A. Create a new Direct Connect gateway in the Company B account. Associate the Company B transit gateway with the new Direct Connect gateway. Create a transit VIF on the existing hosted connection for Company B.

B. Create an association proposal from the Company B account to associate the Company B transit gateway with the Company A Direct Connect gateway. Accept the transit gateway association proposal by logging into the Company A account.

C. Create multiple virtual private gateways. Attach the virtual private gateways to each of Company B's application VPCs. Create a hosted private VIF for each virtual private gateway.

D. Create a new Direct Connect gateway in the Company B account. Associate the Company B transit gateway with the new Direct Connect gateway. Create a hosted private VIF for Company B.

Answer: B

Explanation:

Question: 234

A company has developed a web service for language translation. The web service's application runs on a fleet of Amazon EC2 instances that are in an Auto Scaling group. The instances run behind an Application Load Balancer (ALB) and are deployed in a private subnet. The web service can process requests that contain hundreds of megabytes of data.

The company needs to give some customers the ability to access the web service. Each customer has its own AWS account. The company must make the web service accessible to approved customers without making the web service accessible to all customers.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)

A. Create VPC peering connections with the approved customers only.

B. Create an AWS PrivateLink endpoint service. Configure the endpoint service to require acceptance that will be granted to approved customers only.

C. Configure an authentication action for the endpoint service's load balancer to allow customers to log in by using their AWS credentials. Provide only approved customers with the URL.

D. Configure a Network Load Balancer (NLB) and a listener with the ALB as a target. Associate the NLB with the endpoint service.

E. Associate the ALB with the endpoint service.

Answer: B, D

Explanation:

Question: 235

A company is migrating an application to the AWS Cloud. The company has successfully provisioned and tested connectivity between AWS Direct Connect and the company's on-premises data center. The application runs on Amazon EC2 instances across multiple Availability Zones. The instances are in an Auto Scaling group.

The application communicates through HTTPS to a third-party vendor's data service that is hosted at the company's data center. The data service implements a static ACL through explicit allow listing of client IP addresses.

A network engineer must design a network solution so that the migrated application can continue to access the vendor's data service as the application scales.

Which solution will meet these requirements with the LEAST amount of ongoing change to the vendor's allow list?

A. Configure a private NAT gateway in the subnets for each Availability Zone that the application runs in. Configure the application to target the NAT gateways instead of the data service directly. Update the data service's allow list to include the IP addresses of the NAT gateways.

B. Configure an elastic network interface in the subnets for each Availability Zone that the application runs in. Associate the elastic network interfaces with the Auto Scaling group for the application. Update the data service's allow list to include the IP addresses of the elastic network interfaces.

C. Configure an elastic network interface in the subnets for each Availability Zone that the application runs in. Launch an EC2 instance into each subnet. Attach the respective elastic network interfaces to the new EC2 instances. In the application subnet route tables, configure the new EC2 instances as the next destination for the data service. Update the data service's allow list to include the IP addresses of the elastic network interfaces.

D. Configure an Application Load Balancer (ALB) in the subnets for each Availability Zone that the application runs in. Configure an ALB-associated target group that contains a target that uses the IP address for the data service. Configure the application to target the ALB instead of the data service directly. Update the data service's allow list to include the IP addresses of the ALBs.

Answer: A

Explanation:

Question: 236

A company has a highly available application that is hosted in multiple VPCs and in two on-premises data centers. All the VPCs reside in the same AWS Region. All the VPCs require access to each other and to the on-premises data centers for the transfer of

files that are multiple gigabytes in size.

A network engineer is designing an AWS Direct Connect solution to connect the on-premises data centers to each VPC.

Which architecture will meet the company's requirements with the LEAST operational overhead?

A. Configure a virtual private gateway and a private VIF in each VPC in the Region. Configure a Direct Connect gateway. Associate the VIF of every VPC with the Direct Connect gateway. Create a new private VIF that connects the Direct Connect gateway to each on-premises data center. Configure the new private VIF to exchange BGP routes with the on-premises data centers and to have an MTU of 9001. Use VPC peering between each VPC. Configure static routing in each VPC to provide inter-VPC routing.

B. Configure a virtual private gateway and a private VIF in each VPC in the Region. Configure a Direct Connect gateway. Associate the VIF of every VPC with the Direct Connect gateway. Create a new private VIF that connects the Direct Connect gateway to each on-premises data center. Configure the new private VIF to exchange BGP routes with the on-premises data centers and to have an MTU of 8500. Use VPC peering between each VPC. Configure static routing in each VPC to provide inter-VPC routing.

C. Configure a transit gateway in the same Region of each VPC. Attach each VPC to the transit gateway. Configure a Direct Connect gateway. Associate the Direct Connect gateway with the transit gateway. Associate a new transit VIF with each Direct Connect connection. Configure the new transit VIF to exchange BGP routes and to have an MTU of 9001. Configure route propagation between each VPC and the transit gateway.

D. Configure a transit gateway in the same Region of each VPC. Attach each VPC to the transit gateway. Configure a Direct Connect gateway. Associate the Direct Connect gateway with the transit gateway. Associate a new transit VIF with each Direct Connect connection. Configure the new transit VIF to exchange BGP routes and to have an MTU of 8500. Configure route propagation between each VPC and the transit gateway.

Answer: D

Explanation:

Question: 237

A company has deployed an application in which the front end of the application communicates with the backend instances through a Network Load Balancer (NLB) in the same VPC. The application is highly available across two Availability Zones. The company wants to limit the amount of traffic that travels across the Availability Zones. Traffic from the front end of the application must stay in the same Availability Zone unless there is no healthy target in that Availability Zone behind the NLB. If there is no healthy target in the same Availability Zone, traffic must be sent to the other Availability Zone.

Which solution will meet these requirements?

A. Create a private hosted zone with weighted routing for each Availability Zone. Point the primary record to the local Availability Zone NLB DNS record. Point the secondary record to the Regional NLB DNS record. Configure the front end of the application to perform DNS lookups on the local private hosted zone records.

B. Turn off cross-zone load balancing on the NLB. Configure the front end of the application to perform DNS lookups on the local Availability Zone NLB DNS record.

C. Create a private hosted zone. Create a failover record for each Availability Zone. For each failover record, point the primary record to the local Availability Zone NLB DNS record and point the secondary record to the Regional NLB DNS record. Configure the front end of the application to perform DNS lookups on the local private hosted zone records.

D. Enable sticky sessions (session affinity) so that the NLB can bind a user's session to targets in the same Availability Zone.

Answer: B

Explanation:

Question: 238

A company needs to protect against potential botnet command and control traffic from any Amazon EC2 instances that is in in the company's AWS Environment.

Which solution will meet these requirements?

A. Use AWS Shield Advanced. Activate Shield Advanced protections on the EC2 instances to filter and block botnet traffic.

B. Use Amazon Route 53 Resolver DNS Firewall. Add a rule to a rule group to use the AWSManagedDomainsBotnetCommandandControl managed domain list with an action to block botnet traffic.

C. Use AWS WAF Bot Control. Configure a managed rule group that uses an AWS managed rule set to block botnet traffic.

D. Use AWS Systems Manager. Run a Systems Manager Automation runbook on the EC2 instances to configure the instances to block botnet traffic.

Answer: B

Explanation:

Question: 239

A company has two on-premises data centers. The first data center is in the us-east-1 Region. The Second data canter is in the us-east-2 Region. Each data center connects to the closest AWS Direct Connect facility. The company uses Direct Connect connections, transit VIFs, and a single Direct Connect gateway to establish connectivity to VPCs in us-east-1 and us-east-2 from the company's data centers. The company also has private connectivity from a telecommunications provider that connects the first data center to the second data center.

Recently, there have been multiple connection disruptions to the private connectivity between the data centers. The company needs a solution to improve the reliability of the connection between the two data centers.

Which solution will meet these requirements?

A. Create a new Direct Connect gateway. Enable the Direct Connect SiteLink feature on the transit VIF. Share the CIDR blocks from the first data center and the second data center with each other.

B. Create a new public VIF to both Regions. Enable the Direct Connect SiteLink feature on the new public VIF.

C. Enable the Direct Connect SiteLink feature on the existing Direct Connect connections.

D. Enable the Direct Connect SiteLink feature on the existing transit VIFS that are attached to the existing Direct Connect gateway.

Answer: D

Explanation:

Question: 240

A company has VPCs across 50 AWS accounts and is using AWS Organizations. The company wants to implement web filtering. The requirements for how the traffic must be filtered are the same for all the VPCs. A network engineer plans to use AWS Network Firewall. The network engineer needs to implement a solution that minimizes the number of firewall policies and rule groups that are necessary for this web filtering.

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

A. Create a firewall policy or rule group in each account.

B. Use SCPs to share the firewall policy or rule group.

C. Create a firewall policy or rule group in the management account

D. Use AWS Resource Access Manager (AWS RAM) to share the firewall policy or rule group.

E. Enable sharing within Organizations.

F. Create OUs to share the firewall policy or rule group.

Answer: C, D, E

Explanation:

Question: 241

A company is replatforming a legacy data processing solution to AWS. The company deploys the solution on Amazon EC2 Instances in private subnets that are in one VPC.

The solution uses Amazon S3 for abject storage. Both the data that the solution processes and the data the solution produces are stored in Amazon S3. The solution uses Amazon DynamoDB to save its own state. The company collects flow logs for the VPC. The solution uses one NAT gateway to register its license through the internet. A software vendor provides a specific hostname so the solution can register its license.

The company notices that the AWS bill exceeds the projected budget for the solution. A network engineer uses AWS Cost Explorer to investigate the bill. The network engineer notices that the USE2- NatGateway-Bytes($) usage type is the root cause of the higher than expected bill.

What should the network engineer do to resolve the issue? (Choose two.)

A. Set up Amazon VPC Traffic Mirroring. Analyze the traffic to identify the traffic that the NAT gateway processes.

B. Examine the VPC flow logs to identity the traffic that traverses the NAT gateway.

C. Set up an AWS Cost and Usage Report in the AWS Billing and Cost Management console. Examine the report to find more details about the NAT gateway charges.

D. Verify that the security groups attached to the EC2 instances allow outgoing traffic only to the IP addresses that the hostname resolves to, the VPC CIDR block, and the AWS IP address ranges for Amazon S3 and DynamoDB.

E. Verify that the gateway VPC endpoints for Amazon S3 and DynamoDB are both set up and associated with the route tables of the private subnets.

Answer: B, E

Explanation:

Question: 242

A company is running its application servers on Amazon EC2 instances. The EC2 instances run in separate VPCs that are connected by a transit gateway. The EC2 instances launch in a private subnet with a route to the transit gateway for internal and external connectivity. The external connectivity is provided by a VPC with firewall devices that perform an inspection for packets that ingress and egress through an internet gateway.

A network engineer needs to help the company's application team increase the payload size per packet delivery between the EC2 instances. All network connectivity must be through the transit gateway

What should the network engineer do to meet these requirements?

A. Enable jumbo frames on the transit gateway. Instruct the application team to set the maximum transmission unit (MTU) of the system's network interfaces to 9001 bytes.

B. Instruct the application team to set the maximum transmission unit (MTU) of the VPC to 8500 bytes.

C. Instruct the application team to set up enhanced networking on the system by using the enhanced networking adapter. Set the maximum transmission unit (MTU) to 9001 bytes.

D. Instruct the application team to set the maximum transmission unit (MTU) of the system's network interfaces to 8500 bytes.

Answer: D

Explanation:

Question: 243

A network engineer needs to monitor internet metrics for an application that is in a VPC. The metrics include user experiences such as health events, latency, and traffic insights.

The network engineer sets up Amazon CloudWatch Internet Monitor for the application. The engineer wants to push the internet health events to a third-party target.

Which solution will meet these requirements with the LEAST implementation effort?

A. Create a third-party API endpoint in Amazon EventBridge. Configure internet Monitor to send the events to the third-party API endpoint in EventBridge.

B. Create a third-party API endpoint in Amazon EventBridge. Create a rule in EventBridge that uses Internet Monitor as the source and the third-party API endpoint in EventBridge as the destination.

C. Create a third-party API endpoint in internet Monitor. Configure Internet Monitor to send the events to an Amazon S3 bucket. Configure an AWS Lambda function to send the events to the third- party API endpoint in Internet Monitor.

D. Create a third-party API endpoint in Internet Monitor. Configure Internet Monitor to send the events to the third-party API endpoint in Internet Monitor.

Answer: D

Explanation:

Question: 244

A company has a web application that runs in eight AWS Regions. In each Region, the application is hosted on multiple compute resources behind an Application Load Balancer (ALB).

The different Regions are using different domains. Each ALB is configured to accept only HTTPS traffic. Each ALB uses a certificate from AWS Certificate Manager (ACM).

The company wants to simplify the application's appearance on the web by using a new single domain for all Regions. A network engineer needs to implement this change by designing a solution that also will minimize latency for the application's end users.

Which combination of actions will meet these requirements? (Choose three.)

A. Use ACM to create an SSL/TLS certificate in the us-east-1 Region for the new domain.

B. Set up latency-based routing in Amazon Route 53 for the new domain. Add the ALBs from all the Regions as targets.

C. Create an alias record for the accelerator in Amazon Route 53 for the new domain.

D. Create a standard accelerator in AWS Global Accelerator. Configure a listener for TCP traffic. Add all the ALBs as targets for the listener.

E. Use ACM to create an SSLITLS certificate for each Region. Configure all the ALBs to use the certificate in their respective Regions.

F. Create a custom routing accelerator in AWS Global Accelerator. Configure a listener for HTTPS traffic. Add all the ALBs as targets for the listener. Configure the accelerator to terminate TLS by using the SSLITLS certificate from ACM.

Answer: A, B, F

Explanation:

Question: 245

A company has a VPC that includes application workloads that run on Amazon EC2 instances in a single AWS Region. The company wants to use AWS Local Zones to deploy an extension of the application workloads that run in the Region. The extended workloads in the Local Zone need to communicate bidirectionally with the workloads in the VPC in the Region.

Which solution will meet these requirements MOST cost-effectively?

A. Create a new VPC in the Local Zone. Attach all the VPCs to a transit gateway. Configure routing for the transit gateway and the VPCs. Deploy instances in the new VPC.

B. Deploy a third-party appliance in a new VPC in the Region. Create a new VPC in the Local Zone. Create VPN connections to the appliance for the VPCs. Deploy instances in the new VPC in the Local Zone.

C. Create a new subnet in the Local Zone. Deploy a third-party appliance in the VPC with interfaces in each subnet. Configure the new subnet to route the Local Zone through the appliance. Deploy instances in the new subnet.

D. Create a new subnet in the Local Zone. Configure the new subnet to use a CIDR block that is within the VPC's CIDR block. Deploy instances in the new subnet in the Local Zone.

Answer: D

Explanation:

Question: 246

A company is using AWS Cloud WAN with one edge location in the us-east-1 Region and one edge location in the us-west-1 Region.

A shared services segment exists at both edge locations. Each shared services segment has a VPC attachment to each inspection VPC in each Region. The inspection VPCs inspect traffic from a WAN by using AWS Network Firewall.

The company creates a new segment for a new business unit (BU) in the us-east-1 edge location. The new BU has three VPCs that are attached to the new BU segment. To comply with regulations, the BU VPCs must not communicate with each other. All internet-bound traffic must be inspected in the inspection VPC.

The company updates VPC route tables so any traffic that is bound for internet goes to the AWS Cloud WAN core network.

The company plans to add more VPCs for the new BU in the future. All future VPCs must comply with regulations.

Which solution will meet these requirements in the MOST operationally efficient way? (Choose two.)

A. Update the network policy to share the shared services segment with the BU segment.

B. Create a network policy to share the inspection service segment with the BU segment.

C. Set the isolate-attachments field to True for the BU segment.

D. Set the isolate-attachments field to False for the BU segment.

E. Update the network policy to add static routes for the BU segment. Configure the shared services segment to route traffic related to VPC CIDR blocks to each respective VPC attachment.

Answer: B, C

Explanation:

Question: 247

A company hosts a highly available, scalable, and resilient application on Amazon EC2 instances that are part of an Auto Scaling group. A network engineer is planning to integrate IPv6 support with the application deployment in phases. The first phase is to enable IPv6 service consumption on the public Network Load Balancers (NLBs) that are deployed across the infrastructure. The target groups for the NLBS are configured as the Auto Scaling groups of the EC2 instances that host the application. The NLBs are configured for dual-stack operation.

During the testing of the first phase, the IPv6 application queries are not reaching the backend servers.

What is the cause of this issue?

A. The subnets where the EC2 instances are deployed do not have IPv6 addresses configured.

B. The route tables for the NLB subnets do not have IPV6 routing configured.

C. The route tables for the EC2 subnets do not have IPV6 routing configured.

D. The security groups that are associated with the NLBs do not allow IPv6 traffic.

Answer: C

Explanation:

Question: 248

A company wants to implement a distributed architecture on AWS that uses a Gateway Load Balancer (GWLB) and GWLB endpoints.

The company has chosen a hub-and-spoke model. The model includes a GWLB and virtual appliances that are deployed into a centralized appliance VPC and GWLB endpoints. The model also includes internet gateways that are configured in spoke VPCs.

Which sequence of traffic flow to the internet from the spoke VPC is correct?

A.

1. An application in a spoke VPC sends traffic to the GWLB endpoint based on the VPC route table configuration.
2. Traffic is delivered securely and privately to the GWLB.
3. The GWLB sends the traffic to a virtual appliance for inspection.
4. Return traffic flows back to the GWLB endpoint and out to the internet through the internet gateway.

B.

1. An application in a spoke VPC sends traffic to the GWLB endpoint based on the VPC route table configuration.
2. Traffic is delivered securely and privately to the GWLB endpoint.
3. The GWLB sets the X-Forwarded-For request header and sends the traffic to a virtual appliance for inspection.
4. Return traffic flows back to the GWLB and out to the internet through an internet gateway.

C.

1. An application in a spoke VPC sends traffic to the GWLB endpoint.
2. Traffic is delivered securely and privately to the GWLB.
3. The GWLB sets the X-Forwarded-For request header and sends the traffic to a virtual appliance for inspection.
4. Return traffic flows back to the GWLB endpoint and out to the internet through the internet gateway.

D.

1. An application in a spoke VPC sends traffic to the GWLB.
2. Traffic is delivered securely and privately to the GWLB endpoint.
3. The GWLB sends the traffic to a virtual appliance for inspection.
4. Return traffic flows back to the GWLB and out to the internet through an internet gateway.

Answer: A

Explanation:

Question: 249

A network engineer needs to provide a list of IP addresses that are sending traffic to an Amazon EC2 instance. VPC flow logs are enabled. The EC2 instance has a single network interface and two assigned IP addresses. However, the flow logs are logging traffic only for the primary IP address. The network engineer needs to determine whether any traffic is being sent to the second IP address of the EC2 instance.

What should the network engineer do to locate the traffic flow for the second IP address?

A. Create a new flow log that includes the pkt-dstaddr field to capture the original destination IP address of the traffic.
B. Create a new flow log that includes the dstaddr field to capture the original destination IP address of the traffic.
C. Create a new flow log that includes the pkt-srcaddr field to capture the original destination IP address of the traffic.
D. Create a new flow log that includes the srcaddr field to capture the original destination IP address of the traffic.

Answer: A

Explanation:

Question: 250

A company has configured an AWS Cloud WAN core network with edge locations in the us-east-1 Region and the us-west-1 Region. Each edge location has two segments: development and staging. The segments use the default core network policy.

The company has attached VPCs to the core network. A development VPC is attached to the development segment in us-east-1 and is configured to use the 10.0.0.0/16 CIDR block. A staging VPC is attached to the staging segment in us-west-1 and is configured to use the 10.5.0.0/16 CIDR block. The company has updated the route tables for both VPCs with a route that directs any traffic for 0.0.0.0/0 to the core network.

The company's network team needs to establish communication between the two VPCs by using the AWS Cloud WAN core network. The network team is not receiving a response during tests of communication between the VPCs. The network team has

verified that security groups and network ACLs are not blocking the traffic.

What should the network team do to establish this communication?

A. Update both VPC route tables to have a new static route. Configure a route on the development VPC to direct the traffic for 10.0.0.0/16 to the development VPC attachment. Configure a route on the staging VPC to direct the traffic for 10.5.0.0/16 to the staging VPC attachment.

B. Update the segment filter to allow traffic on the development and staging segments.

C. Set the isolate-attachments parameter to False for the development and staging segments.

D. Update the core network policy to add a static route for each segment. Configure a route to direct the traffic for 10.0.0.0/16 to the development VPC attachment. Configure a route to direct the traffic for 10.5.0.0/16 to the staging VPC attachment.

Answer: D

Explanation:

Question: 251

A company has VPCs in the us-east-1 Region that are connected to each other through a transit gateway. A network engineer needs to establish an AWS Direct Connect connection between the company's on-premises data center and the transit gateway for the migration of a workload.

The Direct Connect connection is UP according to the ConnectionState metric in Amazon CloudWatch. However, the VIF is DOWN. The network engineer has verified the transit VIF and BGP configurations on the on-premises router and has found no issues. However, the network engineer is unable to ping the Amazon peer IP address.

Which combination of steps should the network engineer take to troubleshoot this issue? (Choose three.)

A. Verify that the correct IP address and subnet mask are in use for the subinterface on the router.

B. Ensure that VLAN trunking is disabled on the router.

C. Verify that the router has a MAC address entry from the AWS endpoint in the Address Resolution Protocol (ARP) table.

D. Verify that the optical signal that is received over the cross connect is optimal.

E. Ensure that the correct VLAN tag is applied on the subinterface configuration on the router.

F. Ensure that TCP port 179 is not being blocked at the on-premises router.

Answer: A, C, E

Explanation:

The IP address and subnet mask configuration on the router's subinterface must match the configuration provided by AWS for the

virtual interface (VIF). A mismatch will prevent successful communication between the on-premises router and AWS.

A MAC address entry in the ARP table is essential for Layer 2 connectivity. If the AWS endpoint's MAC address is not visible in the ARP table, it could indicate a connectivity or VLAN tagging issue.

The correct VLAN tag must be configured for the subinterface because AWS Direct Connect requires VLAN tagging to separate traffic between different virtual interfaces.

Question: 252

A logistics company has multiple VPCs in an AWS Region. The company uses a transit gateway to connect the VPCs. The company has several on-premises offices that connect to the transit gateway by using AWS Site-to-Site VPN connections over the internet. The company has configured one transit gateway VPN attachment for each office.

Route propagation is enabled on all route tables. Each Site-to-Site VPN connection uses two tunnels in an active-passive configuration. The company configured each office with appropriate static routes on both the Site-to-Site VPN connection and the office's customer gateway.

The company wants to use both IPsec tunnels of every office to maximize the overall VPN connection bandwidth.

Which design changes are necessary to meet these requirements?

A. Create an AWS Transit Gateway Connect attachment for each office Use the existing VPN attachments as the transport for the new Connect attachments. Set up a Generic Routing

Encapsulation (GRE) tunnel on each customer gateway that terminates on the Connect attachment for each office. Move the static routes from the transit gateway VPN attachment to the customer gateway for the transit gateway Connect attachment.

B. Enable equal-cost multi-path (ECMP) routing on the transit gateway. Ensure ECMP is supported by and enabled on the customer gateways. Enable ECMP on the Site-to-Site VPN connection. Ensure static routes on the customer gateways have equal metrics and administrative distance.

C. Enable equal-cost multi-path (ECMP) routing on the transit gateway. (Ensure ECMP is supported by and enabled on the customer gateways. Change the routing configuration between the transit gateway and the customer gateways from static routing to BGP. Remove related static routes from the customer gateways.

D. Enable equal-cost multi-path (ECMP) routing on the transit gateway. Ensure ECMP is supported by and enabled on the customer gateways. Change the routing configuration between the transit gateway and the customer gateways from static routing to BGP. Ensure the customer gateway applies the correct community strings to give the transit gateway the ability to perform ECMP forwarding.

Answer: C

Explanation:

To use both IPsec tunnels for maximizing bandwidth, equal-cost multi-path (ECMP) routing must be enabled. ECMP allows the transit gateway to load balance traffic across multiple paths (in this case, both IPsec tunnels). For ECMP to work:

Transit Gateway ECMP Support: The transit gateway must have ECMP routing enabled to distribute traffic across multiple VPN tunnels.

BGP Configuration: Static routing cannot support ECMP. Switching to BGP allows dynamic route advertisements and supports ECMP. Removing static routes ensures that the BGP-learned routes take precedence.

Customer Gateway ECMP Support: The customer gateway must also support ECMP for the configuration to work end-to-end.

By implementing these changes, both tunnels can be utilized simultaneously, effectively increasing the available bandwidth for the Site-to-Site VPN connections.

Question: 253

A finance company runs multiple applications on Amazon EC2 instances in two VPCs that are within a single AWS Region. The company uses one VPC for stock trading applications. The company uses the second VPC for financial applications. Both VPCs are connected to a transit gateway that is configured as a multicast router.

In the stock trading VPC, an EC2 instance that has an IP address of 10.128.10.2 sends trading data over a multicast network to the 239.10.10.10 IP address on UDP Port 5102. The company recently launched two new EC2 instances in the financial application VPC. The new EC2 instances need to receive the multicast stock trading data from the EC2 instance that is in the stock trading VPC.

Which combination of steps should the company take to meet this requirement? (Choose three.)

A. Add the elastic network interfaces of the two new EC2 instances as members of the multicast group by using the group IP address of 239.10.10.10.

B. Add an inbound rule to the security groups that are attached to the multicast receiver instances. Configure the rule as follows:

Protocol: IGMP Version 2. Port: 5102, and Source: 239 10.10.10/32

C. Create associations to two EC2 instance IDs on the financial application VPC transit gateway attachment under the transit gateway multicast domain.

D. Create an association to EC2 instance subnets on the financial application VPC transit gateway attachment under the transit gateway multicast domain.

Add an inbound rule to the security groups that are attached to the multicast receiver instances. Configure the rule as follows.

E. Protocol: UDP, Port: 5102, and Source: 10.128.10.2/32

F. Add an inbound rule to the security groups that are attached to the multicast receiver instances. Configure the rule as follows:

Protocol: IGMP Version 2. Port: All, and Source: 0 0.0.0/32

Answer: A, C, E

Explanation:

Add ENIs to the multicast group: To receive multicast traffic, the ENIs of the receiver EC2 instances must be explicitly added as members of the multicast group using the multicast group IP address (239.10.10.10).

Create associations to the EC2 instance IDs: Multicast domains in a transit gateway allow multicast traffic to flow between VPCs. To enable specific instances to receive the multicast traffic, the instance IDs in the financial application VPC must be associated with the transit gateway multicast domain.

Add a security group rule for UDP traffic: The receiver instances need a security group rule to allow inbound UDP traffic on port 5102 from the sender EC2 instance (source IP: 10.128.10.2). This ensures that multicast traffic is allowed to reach the receiving instances.

Question: 254

A company runs workloads in multiple VPCs in the us-east-1 Region. The VPCs are connected to a transit gateway. An AWS Direct Connect connection provides private connectivity between a data center that is in the US and the transit gateway. A Direct Connect

gateway is associated with the transit gateway.

The company has recently opened a new office location in London. The company plans to launch cloud services in multiple VPCs in the eu-west-2 Region. Users in the new London office must have private access to the workloads that run in us-east-1. Users in the US data center must have access to any workloads that are created in eu-west-2. A network engineer must implement a flexible solution that provides users the required access. The solution must be able to accommodate future growth.

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

A. Create an AWS Site-to-Site VPN connection from the London office to the Direct Connect gateway in us-east-1.

B. Establish a new Direct Connect connection for the London office. Attach the new Direct Connect connection to the existing Direct Connect gateway. Create a transit gateway in eu-west-2. Associate the new transit gateway with the existing Direct Connect gateway. Create a peering connection between the transit gateways in us-east-1 and eu-west-2.

C. Create an AWS Site-to-Site VPN connection from the London office to each of the VPCs that are in us-east-1.

D. Establish a new AWS Direct Connect connection for the London office Create a new Direct Connect gateway and a transit gateway in eu-west-2. Attach the new Direct Connect connection to the new Direct Connect gateway. Create a peering connection between the transit gateways in us-east-1 and eu-west-2.

_____Answer: B

Explanation:

This solution provides a scalable and flexible architecture that accommodates future growth while ensuring connectivity between regions and locations.

New Direct Connect Connection in London: Establishing a Direct Connect connection for the London office provides a high-performance, private, and reliable connection between the office and AWS.

Use the Existing Direct Connect Gateway: By attaching the new Direct Connect connection to the existing Direct Connect gateway, the solution leverages the existing infrastructure, reducing operational complexity.

Transit Gateway in eu-west-2: The transit gateway in eu-west-2 enables centralized routing and connectivity management for multiple VPCs in that Region.

Transit Gateway Peering: Peering the transit gateways in us-east-1 and eu-west-2 ensures seamless private connectivity between the regions, allowing users in the US and London to access workloads in both regions.

This design ensures low operational effort because it avoids creating and managing multiple individual VPN connections. It also supports future growth by allowing additional VPCs, regions, or office locations to connect with minimal configuration changes.

Question: 255

A company has 10 Amazon EC2 instances that run web server software in a production VPC. The company also has 10 web servers that run in an on-premises data center. The company has a 10 Gbps AWS Direct Connect connection between the on-premises data center and the production VPC. The data center uses the 10.100.0.0/20 CIDR block.

The company needs to implement a load balancing solution that receives HTTPS traffic from thousands of external users. The solution must distribute the traffic across the web servers on AWS and the web servers in the data center. Regardless of the location of the web servers, HTTPS requests must go to the same web server for the duration of the session.

Which solution will meet these requirements?

A. Deploy a Network Load Balancer (NLB) in the production VPC. Create one target group for the EC2 Instances and a second target group for the on-premises servers. Specify IP as the target type. Register the EC2 instances and the on-premises servers with the target groups. Enable connection draining on the NLB.

B. Deploy an Application Load Balancer (ALB) in the production VPC. Create one target group for the EC2 Instances and a second target group for the on-premises servers. Specify IP as the target type. Register the EC2 instances and the on-premises servers with the target groups. Enable applicationbased sticky sessions on the ALB.

C. Deploy a Network Load Balancer (NLB) in the production VPC. Create one target group for the EC2 Instances and a second target group for the on-premises servers. Specify instance as the target type. Register the EC2 instances and the on-premises servers with the target groups. Enable sticky sessions on the NLB.

D. Deploy an Application Load Balancer (ALB) in the production VPC. Create one target group for the EC2 Instances and a second target group for the on-premises servers. Specify instance as the target type. Register the EC2 instances and the on-premises servers with the target groups. Enable application-based sticky sessions on the ALB.

Answer: B

Explanation:

Application Load Balancer (ALB) is the best choice for handling HTTPS traffic because it operates at the application layer (Layer 7). It supports features such as sticky sessions, which are required to ensure that requests from the same user are routed to the same server for the duration of the session.

IP as the Target Type: To load balance traffic between both EC2 instances and on-premises servers, the target type must be set to IP. This configuration allows the ALB to distribute traffic to servers identified by IP addresses, including those in the on-premises data center.

Sticky Sessions: Application-based sticky sessions ensure that the load balancer uses cookies to route requests consistently to the same web server, meeting the requirement for session persistence.

Question: 256

A global company is establishing network connections between the company's primary and secondary data centers and a VPC. A network engineer needs to maximize resiliency and fault tolerance for the connections. The network bandwidth must be greater than 10 Gbps.

Which solution will meet these requirements MOST cost-effectively?

A. Set up a 100 Gbps connection at the primary data center that terminates at an AWS Direct Connect location. Set up a second 100 Gbps connection at the secondary data center that terminates at a second Direct Connect location. Ensure the connections are managed by separate providers.

B. Set up a 10 Gbps connection at the primary data center that terminates at an AWS Direct Connect location. Set up a second 10 Gbps connection at the secondary data center that terminates at a second Direct Connect location. Ensure the connections are managed by separate providers.

C. Set up two 10 Gbps connections at the primary data center that terminate at one AWS Direct Connect location. Ensure the connections are managed by separate providers. Set up two 10 Gbps connections at the secondary data center that terminate at a second Direct Connect location. Ensure the connections are managed by separate providers.

D. Set up a 10 Gbps connection at the primary data center that terminates at an AWS Direct Connect location. Set up an AWS Site-to-Site VPN connection at the secondary data center that terminates at a virtual private gateway in the same Region as the company's VPC.

Answer: C

Explanation:

Multiple 10 Gbps Connections: By setting up two 10 Gbps connections at each data center, the solution achieves an aggregate bandwidth of 20 Gbps per data center, exceeding the requirement of 10 Gbps. Using multiple 10 Gbps links is more cost-effective than deploying 100 Gbps links.

Separate Providers: Ensuring that the connections are managed by separate providers minimizes the risk of a single provider's failure affecting the network.

Two Direct Connect Locations: Terminating connections at two Direct Connect locations ensures geographic redundancy. This setup minimizes the impact of outages or disruptions at a single Direct Connect location.

Question: 257

A company's data center is connected to a single AWS Region by an AWS Direct Connect dedicated connection. The company

has a single VPC in the Region. The company stores logs for all its applications locally in the data center.

The company must keep all application logs for 7 years. The company decides to copy all application logs to an Amazon S3 bucket.

Which solution will meet these requirements?

A. Create a public VIF on the Direct Connect connection. Create an Amazon S3 gateway endpoint in the VPC.

B. Create a private VIF on the Direct Connect connection. Create an Amazon S3 gateway endpoint in the VPC.

C. Create a private VIF on the Direct Connect connection. Create an Amazon S3 interface endpoint in the VPC.

D. Create a public VIF on the Direct Connect connection. Create an Amazon S3 interface endpoint in the VPC.

Answer: B

Explanation:

Private VIF on Direct Connect: A private virtual interface (VIF) is used to establish a private connection between your data center and your VPC over AWS Direct Connect. This allows traffic to remain within the private network.

Amazon S3 Gateway Endpoint: A gateway endpoint provides a scalable and highly available way to privately access Amazon S3 without requiring an internet gateway or NAT gateway. Gateway endpoints are more cost-effective than interface endpoints for accessing S3, especially for high- volume data transfers like application logs.

Question: 258

A company is planning to host a secure web application across multiple Amazon EC2 instances. The application will have an associated DNS domain in an Amazon Route 53 hosted zone.

The company wants to protect the domain from DNS poisoning attacks. The company also wants to allow web browsers to authenticate into the application by using a trusted third party.

Which combination of actions will meet these requirements?

A. Configure the Route 53 hosted zone to use DNS Security Extensions (DNSSEC). Install self-signed X.509 certificates on the EC2 instances.

B. Configure a Name Authority Pointer (NAPTR) record in the Route 53 hosted zone. Install X 509 certificates that are signed by a public certificate authority on the EC2 instances.

C. Configure the Route 53 hosted zone to use DNS Security Extensions (DNSSEC). Install X.509 certificates that are signed by a public certificate authority on the EC2 instances.

D. Configure a Name Authority Pointer (NAPTR) record in the Route 53 hosted zone. Install self-signed X.509 certificates on the EC2 instances.

Answer: C

Explanation:

DNSSEC protects against DNS poisoning attacks by enabling authentication of DNS data integrity and origin. When DNSSEC is enabled in the Route 53 hosted zone, it signs the DNS records, ensuring their authenticity.

To enable web browsers to authenticate the application securely using a trusted third party, X.509 certificates signed by a public Certificate Authority (CA) are required. These certificates allow HTTPS communication, ensuring that the web browsers trust the application.

Question: 259

A company is planning to use an AWS Transit Gateway hub and spoke architecture to migrate to AWS. The current on-premises multi-protocol label switching (MPLS) network has strict controls that enforce network segmentation by using MPLS VPNs. The company has provisioned two 10 Gbps AWS Direct Connect connections to provide resilient, high-speed, low-latency connectivity to AWS.

A security engineer needs to apply the concept of network segmentation to the AWS environment to ensure that virtual routing and forwarding (VRF) is logically separated for each of the company's software development environments. The number of MPLS VPNs will increase in the future. Onpremises MPLS VPNs will have overlapping address space. The company's AWS network design must support overlapping address space for the VPNs.

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

A. Deploy a software-defined WAN (SD-WAN) head-end virtual appliance and an SD-WAN controller into a Transit Gateway Connect VPC. Configure the company's edge routers to be managed by the new SD-WAN controller and to use SD-WAN to segment the traffic into the defined segments for each of the company's development environments.

B. Configure IPsec VPNs on the company edge routers for each MPLS VPN for each of the company's development environments. Attach each IPsec VPN tunnel to a discrete MPLS VPN. Configure AWS Site-to-Site VPN connections that terminate at a transit gateway for each MPLS VPN. Configure a transit gateway route table that matches the MPLS VPN for each Transit Gateway VPN attachment.

C. Create a transit VPC that terminates at the AWS Site-to-Site VRF-aware IPsec VPN. Configure IPsec VPN connections to each VPC for each of the company's development environment VRFs.

D. Configure a Transit Gateway Connect attachment for each MPLS VPN between the company's edge routers and Transit

Gateway. Configure a transit gateway route table that matches the MPLS VPN for each of the company's development environments.

Answer: D

Explanation:

Transit Gateway Connect Attachments: Transit Gateway Connect attachments support dynamic routing via BGP and can segment traffic using VRFs, which aligns with the company's requirement for logical network segmentation similar to MPLS VPNs. Each Connect attachment can represent a VRF from the on-premises MPLS network.

Support for Overlapping Address Space: Transit Gateway Connect attachments allow the segmentation of routing through the use of separate transit gateway route tables for each development environment. This ensures that overlapping IP address spaces between the environments remain isolated.

Low Operational Overhead leverages Transit Gateway's native capabilities for network segmentation and routing, minimizing the need for additional software-defined networking solutions or complex VPN configurations.

Question: 260

A company is planning to migrate to AWS and use multiple VPCs in multiple AWS Regions. A network engineer must connect the eu-west-1 and eu-central-1 Regions to the company headquarters and branch office, respectively.

The network engineer created a production VPC, named Prod A, with a CIDR block of 10.0.0.0/16. Prod A runs in an account in eu-west-1. The network engineer then created another production VPC, named Prod B, with a CIDR block of 10.1.0.0/16. Prod B runs in a different account in eu-central-1.

The network engineer performed the following steps to try to achieve the required connectivity:

1. Created one transit gateway in each Region
2. Shared and accepted the transit gateways with the production accounts in both Regions
3. Configured the peering attachment between both transit gateways
4. Attached both VPCs to the respective Region transit gateway
5. Created both transit gateway route tables and associated the attachments with the route tables
6. Configured a static route in both transit gateway route tables to send traffic to the remote VPC in the other Region

7. Activated route propagation on the VPC route tables in each Region

After the configuration, the network engineer tried to connect from Prod A to Prod B. However, the connection was unsuccessful.

What should the network engineer do to achieve the required connectivity?

A. Modify the IP address of the peering attachment to a wider range.

B. Delete the static routes that were in the transit gateway route table to send traffic to the remote VPC and enable route propagation instead.

C. Create a new route destined to 10.0.0.0/8 in both production VPC route tables with the Region transit gateway as the target.

D. Modify the transit gateway route tables from the production accounts to propagate routes dynamically between the production VPCs.

Answer: C

Explanation:

Question: 261

A company hosts an application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are part of an Amazon EC2 Auto Scaling group.

To comply with new security standards, the company must capture all application access data, including server response codes, request paths, latency, and client IP addresses. The company also needs to query the captured data for performance analysis.

Which solution will meet these requirements?

A. Enable VPC flow logs on the ALB subnets. Store the logs to an Amazon S3 bucket. Query the logs in the S3 bucket by using Amazon Athena.

B. Configure Amazon VPC Traffic Mirroring on all EC2 elastic network interfaces. Deploy a third-party monitoring appliance from AWS Marketplace in a private subnet. Use Amazon Data Firehose to send all mirrored traffic to the monitoring appliance. Query the logs directly from the monitoring appliance.

C. Configure Amazon CloudWatch detailed monitoring on the EC2 instances Include all available logs. Use Amazon Data Firehose to send all the collected logs to an Amazon S3 bucket. Query the data directly from the S3 bucket.

D. Enable access logs on the ALB. Store the logs in an Amazon S3 bucket. Query the logs in the S3 bucket by using Amazon Athena.

Answer: D

Explanation:

ALB Access Logs: Enabling access logs on the ALB provides detailed information about incoming requests, including server response codes, request paths, latency, and client IP addresses. These logs are specifically designed to capture the required application access data.

S3 for Storage: The logs are stored in Amazon S3, providing a cost-effective and scalable solution for retaining the data.

Athena for Querying: Amazon Athena can be used to query the data directly from the S3 bucket without the need to move or transform the data, making it efficient for performance analysis.

Question: 262

A company has five VPCs in the us-east-1 Region. The company hosts an internal web application in us-east-1. One of the company's VPCs. named VPC-A, needs to connect to an external partner's AWS environment. The partner's environment is in the same AWS Region where the partner hosts a new version of the company's web application. The partner hosts its version of the application in a VPC named VPC-B.

The company has Amazon EC2 instances in VPC-A that need to connect to the web application in VPC-B A network engineer notices that the partner's VPC-B and the company's VPC-A use the same IP space. The network engineer needs a solution to allow the EC2 instances to connect to the web application. The solution must not negatively affect the exiting environment of the company or the partner.

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

A. Establish a VPC peering connection between VPC-A to VPC-B.

B. Ensure the partner creates a VPC endpoint service that uses a Network Load Balancer in VPC-B.

C. Deploy a VPC endpoint in VPC-A that uses a VPC endpoint service that is shared by the partner.

D. Deploy a new routable VPC CIDR block as a secondary CIDR block to both VPC-A and VPC-B. Deploy a public NAT gateway

in VPC-A.

E. Establish an AWS Site-to-Site VPN connection between VPC-A and VPC-B.

Answer: B, C

Explanation:

The partner can create a VPC endpoint service backed by an NLB in VPC-B. The NLB routes traffic to the web application in VPC-B, allowing external access to the application without requiring direct VPC-to-VPC connectivity. This approach ensures isolation and avoids conflicts due to overlapping IP spaces.

The company can deploy a VPC endpoint in VPC-A to connect to the VPC endpoint service in VPC-B. This setup allows EC2 instances in VPC-A to securely connect to the web application in VPC-B without modifying existing routing or NAT configurations.

Question: 263

A company has a hybrid environment that connects an on-premises data center to the AWS Cloud. The hybrid environment uses a 10 Gbps AWS Direct Connect dedicated connection. The Direct Connect connection has multiple private VIFs that terminate in multiple VPCs.

To comply with regulations, the company must encrypt all WAN traffic, regardless of the underlying transport. The company needs to implement an encryption solution that will not affect the company's bandwidth capacity.

Which solution will meet these requirements?

A. Create a public VIF. Configure a new AWS Site-to-Site VPN connection to use the new public VIF.

B. Configure MAC security (MACsec) support on the port of the existing Direct Connect connection. Change the encryption mode to must_encrypt.

C. Configure a new Direct Connect connection that supports MAC security (MACSec) Associate the existing VIFs to the new Direct Connect connection.

D. Create a public VIF. Configure a new private IP VPN that uses the Direct Connect connection.

Answer: B

Explanation:

MACsec for Direct Connect:

MACsec (Media Access Control Security) is an IEEE standard (802.1AE) for encrypting traffic at Layer 2. AWS Direct Connect supports MACsec on dedicated connections of 10 Gbps and 100 Gbps capacity. This ensures that all WAN traffic over the Direct Connect connection is encrypted, meeting regulatory requirements.

Does Not Affect Bandwidth:

MACsec operates at the physical layer (Layer 2), and its encryption overhead is negligible. This ensures that the company's bandwidth capacity is not affected.

Existing Direct Connect Connection:

Configuring MACsec on the port of the existing Direct Connect connection avoids the need to establish a new connection, reducing complexity and costs.

Question: 264

A company needs to capture and log traffic for Nitro-based Amazon EC2 instances to comply with regulations. The company's network team has prepared a solution that enables VPC traffic mirroring and sends traffic to a second set of EC2 instances in an Auto Scaling group.

The network team has added a Network Load Balancer (NLB) in front of the EC2 instances the traffic will be sent to. However, the solution does not send any mirrored traffic to the EC2 instances that are behind the NLB.

How should the network team configure traffic mirroring to use the NLB endpoint?

A. Select the NLB as a source for traffic mirroring. Use a UDP listener.

B. Select the NLB as a target for traffic mirroring. Use a TCP listener and a UDP listener.

C. Select the NLB as a target for traffic mirroring. Use a TCP listener.

D. Select the NLB as a target for traffic mirroring. Use a UDP listener.

Answer: D

Explanation:

Traffic Mirroring with UDP: VPC traffic mirroring sends mirrored traffic as UDP packets encapsulated in VXLAN (Virtual eXtensible Local Area Network). Therefore, the target for traffic mirroring must have a UDP listener to accept the mirrored traffic.

NLB as a Target: The Network Load Balancer (NLB) can be configured as the target for traffic mirroring. By setting up a UDP listener on the NLB, it can properly receive the VXLAN-encapsulated mirrored traffic and forward it to the EC2 instances behind it.

Correct Listener Protocol: Since mirrored traffic uses UDP, setting up a TCP listener or using both TCP and UDP listeners will not work. Only a UDP listener can correctly handle the traffic.

Question: 265

A US-based company is expanding its business to Europe. A network engineer needs to extend the company's network infrastructure by setting up a new hub and spoke architecture in the eu-west-1 Region. The network engineer uses a transit gateway peering connection to connect the new resources in eu-west-1 to an existing environment in the us-east-1 Region.

The hub and spoke architecture in each AWS Region includes an inspection VPC that uses AWS Network Firewall to centralize traffic inspection for each Region. To reduce costs, the network engineer decides to inspect inter-Region traffic by using the inspection VPC in the Region that originates the traffic. The network engineer configures the transit gateway route tables accordingly for each Region.

When the network engineer tests the new architecture, communication within each Region works as expected. However, the network engineer finds that inter-Region communication is not working. The network engineer must resolve the inter-Region communication issue.

Which solution will meet this requirement?

A. Configure Open Shortest Path First (OSPF) routing on the transit gateway peering connection to propagate the VPC CIDR blocks from each Region to the remote peer.

B. Use AWS Resource Access Manager (AWS RAM) to share access between the transit gateways. Enable the Allow sharing with anyone setting.

C. Prevent asymmetric routing in the inspection VPCs by ensuring that both requests and responses are inspected by the same inspection VPC

D. Enable Appliance mode on both the transit gateway attachments for the inspection VPC.

Answer: D

Explanation:

Inspection VPC with Network Firewall: When using a central inspection VPC with AWS Network Firewall, traffic must be routed through the inspection VPC for inspection. This requires the traffic to pass through transit gateway attachments that are configured to handle such scenarios.

Appliance Mode: Enabling Appliance mode on the transit gateway attachments for the inspection VPC allows asymmetric traffic flows, where packets in one direction and their corresponding return packets can take different paths. This is crucial for inspection because traffic may be routed back differently after being inspected.

Inter-Region Communication: Without Appliance mode, transit gateways drop traffic if the return path is not symmetric. Enabling Appliance mode ensures that the inspection VPC can handle interRegion traffic flows without breaking communication.

Question: 266

A company runs applications in two VPCs that are in separate AWS Regions. One VPC is in the us- east-1 Region. The second VPC is in the us-west-1 Region. The company needs to establish connectivity between the two VPCs. The company also needs to connect the VPCs to applications that run in an on-premises data center.

The current traffic requirement between the VPCs is 50 TB per month. The company expects traffic volume between the VPCs to increase. The traffic requirement from the VPCs to the on-premises data center is 10 TB per month. The company expects the traffic between the VPCs and the data center to remain constant.

Which solution will meet these requirements MOST cost-effectively?

A. Create a transit gateway in each Region. Create VPN connections from the transit gateways to the on-premises firewall. Create a peering connection between the transit gateways.

B. Create a virtual private gateway in each Region. Create VPN connections from the on-premises firewall to the virtual private gateways. Configure the on-premises firewall to route the traffic between the two VPCs.

C. Create a virtual private gateway in each Region. Create VPN connections from the on-premises firewall to the virtual private gateways. Create a VPC peering connection between the two VPCs.

D. Create a virtual private gateway in each Region. Create VPN connections from the on-premises firewall to the virtual private gateways. Create a VPN connection between the virtual private gateways.

Answer: A

Explanation:

Traffic Volume Consideration: The traffic volume between the VPCs (50 TB per month and increasing) justifies the use of transit gateways, which are designed for scalable, high-throughput

interconnectivity. A VPC peering connection would not scale as efficiently for this traffic volume.

On-Premises Connectivity: Establishing VPN connections from the on-premises firewall to the transit gateways ensures secure connectivity between the on-premises data center and both VPCs.

Transit Gateway Peering: Creating a peering connection between the transit gateways allows for efficient inter-Region communication between the VPCs without routing through the on-premises data center, reducing latency and costs.

Cost Efficiency: Transit gateway peering provides a cost-effective solution for large inter-Region traffic volumes compared to alternatives like routing all traffic through the on-premises data center, which would incur higher egress costs and potentially create a bottleneck.

Question: 267

A company runs workloads in multiple VPCs. The company needs to securely access a workload in one of the VPCs, named VPC-A, from an on-premises data center. A network engineer sets up an AWS Site-to-Site VPN connection to a transit gateway. The network engineer configures dynamic routing for the connection, and communication works properly.

Recently, the owner of VPC-A added another CIDR range to the VPC. The VPC-A owner created workloads that use the additional CIDR range.

The company's on-premises network is unable to reach the new workloads. The network engineer needs to resolve the network connectivity issue and ensure that connectivity will not be affected if additional VPC CIDR ranges are added to the VPC in the future.

Which solution will meet these requirements with the MOST operational efficiency?

A. Configure route propagation for VPC-A to the VPN attachment route table.

B. Manually update the VPN attachment route table to include the new CIDR range.

C. Configure an Amazon EventBridge rule to invoke an AWS Lambda function when the rule to matches an update to the VPC-A CIDR range. Configure the Lambda function to update the VPN attachment route table.

D. Configure an Amazon CloudWatch alarm to invoke an AWS Lambda function when there is an update to the VPC-A CIDR range. Configure the Lambda function to update the VPN attachment route table. Restart the VPN tunnels.

Answer: A

Explanation:

Dynamic Route Propagation: By enabling route propagation on the VPN attachment route table, any changes to the VPC CIDR (such as adding new CIDR ranges) are automatically propagated to the transit gateway route tables and the on-premises network. This ensures seamless connectivity without requiring manual updates.

Operational Efficiency: This approach eliminates the need for manual updates or additional automation, reducing operational overhead. Any new CIDR ranges added to VPC-A will automatically be advertised to the on-premises network through the VPN connection.

Future-Proof Solution: Configuring route propagation ensures that future updates to VPC CIDR ranges are automatically handled, providing a robust and scalable solution.

Question: 268

A company is migrating its internet VPN connections to dedicated AWS Direct Connect connections. The company needs to set up the Direct Connect connections so that all network communications are encrypted in transit.

Which combination of steps will meet this requirement? (Choose three.)

A. Create new Direct Connect connections while requesting MACsec ports.

B. Create a MACsec Connectivity Association Key Name (CKN) and Connectivity Association Key (CAK) pair. Associate the pair with each new connection.

C. Update the on-premises routers to use MACsec and the shared Connectivity Association Key Name (CKN) and Connectivity Association Key (CAK) pair.

D. Create a shared key for an IPsec connection.

E. Configure a new Direct Connect gateway. Associate the shared key with the new Direct Connect gateway.

F. Set up IPsec on the on-premises router. Associate the shared key with the IPsec configuration.

Answer: A, B, C

Explanation:

MACsec (Media Access Control Security) is a Layer 2 encryption standard that can be enabled on AWS Direct Connect dedicated connections. Requesting MACsec ports ensures that the new connections support encryption at the physical layer.

MACsec requires a CKN and CAK pair to establish encrypted communication between the onpremises routers and the AWS Direct Connect routers. This ensures the encryption keys are securely shared and recognized by both endpoints.

To enable MACsec on the Direct Connect link, the on-premises routers must be configured to use the same CKN and CAK. This ensures secure communication and proper encryption over the connection.

Question: 269

A company has an application VPC and a networking VPC that are connected through VPC peering. The networking VPC contains a Network Load Balancer (NLB). The application VPC contains Amazon EC2 instances that run an application. The EC2 instances are part of a target group that is associated with the NLB in the networking VPC.

The company configures a third VPC and peers it to the networking VPC. The new VPC contains a new version of the existing application. The new version of the application runs on new EC2 instances in an application subnet. The new version of the application runs in a different Availability Zone than that original version of the application.

The company needs to establish connectivity between the NLB and the new version of the application.

Which combination of steps will meet this requirement? (Choose three.)

A. Register the new application EC2 instances with the NLB by using the instance IDs.

B. Register the new application EC2 instances with the NLB by using instance IP addresses.

C. Configure the NLB in the Availability Zone where the new application EC2 instances run.

D. Configure the NLB to use zonal shift.

E. Configure the network ACL for the application subnet in the new VPC to allow outbound connections.

F. Configure the network ACL for the application subnet in the new VPC to allow inbound connections and outbound connections.

Answer: B, C, F

Explanation:

Since the new application EC2 instances reside in a different VPC, the NLB must use IP addresses to register the instances. Instance IDs can only be used for targets within the same VPC as the NLB.

To handle traffic efficiently, the NLB must be configured in the Availability Zone where the new application EC2 instances are running. This ensures proper routing and load balancing for instances in that zone.

The network ACL for the new application's subnet must allow both inbound and outbound traffic to facilitate communication between the NLB in the networking VPC and the application EC2 instances in the new VPC.

Question: 270

A company uses AWS Site-to-Site VPN connections to encrypt traffic between the company's onpremises location and a single VPC. The Site-to-Site VPN connections use two 1 Gbps AWS Direct Connect connections with public VIFs. The company plans to add 15 additional VPCs in the same AWS Region.

The company must maintain the same level of encryption that the Site-to-Site VPN connections currently provide for each connection between the on-premises location and the new VPCs. The new connections must not use public IP addresses. The bandwidth of the Site-to-Site VPN connections will remain less than the current provisioned speed.

Which combination of steps will meet these requirements with LEAST operational overhead? (Choose three.)

A. Create a transit gateway and a Direct Connect gateway. Associate the transit gateway with the Direct Connect gateway. Attach all the new VPCs to the transit gateway.

B. For each new VPC, create a new Direct Connect private VIF to a Direct Connect gateway. Associate all VPCs with the Direct Connect gateway.

C. Assign a private IP CIDR block to the transit gateway.

D. Assign a public IP CIDR block to the transit gateway.

E. Create a transit VIF to the Direct Connect gateway. Create a Site-to-Site VPN private IP VPN connection.Create a public VIF.

F. Create a Site-to-Site VPN public IP VPN connection.

Answer: A, C, E

Explanation:

The transit gateway allows for scalable and centralized routing between multiple VPCs and onpremises networks. Associating it with a Direct Connect gateway enables private connectivity over the existing Direct Connect connections, which is more efficient than creating separate Direct Connect VIFs for each new VPC.

Assigning a private IP CIDR block to the transit gateway ensures that all traffic is routed securely and avoids the use of public IPs, which meets the requirement for private connectivity.

A transit VIF to the Direct Connect gateway provides encrypted communication using a private IP VPN. This configuration ensures encryption similar to the Site-to-Site VPN connections while leveraging private IPs and the existing Direct Connect infrastructure.

Question: 271

A company hosts application servers on premises and on Amazon EC2 instances in a VPC. The application servers access data that is hosted in an Amazon S3 bucket through the public internet. The EC2 instances in the VPC use an AWS Site-to-Site VPN for connectivity with the on-premises application servers.

New company regulations state that all traffic between the application servers and the S3 bucket must remain private and must not use public IP addresses.

Which solution will meet these requirements MOST cost-effectively?

A. Configure an S3 gateway endpoint Modify the route table with the appropriate route for the endpoint. Access the S3 bucket through the gateway endpoint from the EC2 instances.

B. Configure an S3 interface endpoint. Update the on-premises servers and EC2 instances to use the interface endpoint DNS name to access the S3 bucket.

C. Configure an S3 interface endpoint. Update the on-premises servers to use the interface endpoint DNS name to access the S3 bucket. Configure an S3 gateway endpoint. Modify the route table so that the EC2 instances use the gateway endpoint.

D. Configure an S3 gateway endpoint. Modify the route table with the appropriate route for the endpoint. Use an S3 bucket policy to restrict access to the gateway endpoint. Configure a proxy server fleet behind a Network Load Balancer in the VPC so that the on-premises servers can access the S3 bucket.

Answer: A

Explanation:

S3 Gateway Endpoint for VPC Access: A gateway endpoint is a highly cost-effective and scalable solution for providing private access

to Amazon S3 from within a VPC. By creating an S3 gateway endpoint and modifying the route table, EC2 instances in the VPC can access S3 without using public IPs or traversing the internet.

No Need for Interface Endpoints or Additional Proxies: Interface endpoints are more expensive and generally not required for accessing S3 unless you need fine-grained network traffic control or private DNS resolution. Similarly, deploying a proxy fleet introduces unnecessary complexity and costs.

Simplicity and Cost-Effectiveness: The S3 gateway endpoint meets the requirement of keeping traffic private and is the simplest, most cost-effective solution for accessing S3 from the EC2 instances in the VPC.

Question: 272

A company uses AWS Network Firewall to protect outgoing traffic for multiple VPCs that are in the same AWS account. Each VPC contains Amazon EC2 instances that host the company's applications. Each EC2 instance is tagged with the name of the application it hosts. The EC2 instances are in Auto Scaling groups.

A Network Firewall stateful rule group must remain up-to-date, even when an Auto Scaling group launches and terminates EC2 instances.

Which solution will meet this requirement with the LEAST implementation and administrative effort?

A. Create a network ACL for each application. Reference the network ACL in the stateful rule group.

B. Create a prefix list for each application. Reference the prefix list in the stateful rule group.

C. Create an AWS Lambda function that queries the EC2 instance tags for each application name and then updates the stateful rule group with the IP address of each instance.

D. Create a resource group for each application name. Reference the Amazon Resource Name (ARN) for the resource groups in the stateful rule group.

Answer: B

Explanation:

Prefix Lists: AWS prefix lists are a simple way to manage IP address ranges that can be reused across security groups, route tables, and firewall rules. By creating a prefix list for each application, you can group the IP addresses of the EC2 instances that host a specific application.

Dynamic Updates with Minimal Effort: When EC2 instances are launched or terminated in Auto Scaling groups, the IP ranges can be dynamically updated in the prefix lists. This approach avoids manual updates to the stateful rule group while ensuring that the rule group remains up-to-date.

Low Administrative Overhead: Referencing the prefix lists in the stateful rule group minimizes the administrative effort, as any changes to the prefix lists automatically reflect in the firewall rules without needing direct modifications to the rule group.

Question: 273

A company has multiple AWS Site-to-Site VPN connections between an on-premises environment and multiple VPCs. The Site-to-Site VPN connections use virtual private gateways and are configured with IPv4 addresses. The company hosts several internal applications in the VPCs.

Application users have reported that the applications are performing slowly. A network engineer notices excessive latency in the network path that the VPN connections use. The network engineer needs to resolve the excessive latency.

Which solution will meet this requirement?

A. Use AWS Global Accelerator to deploy an accelerator on the existing Site-to-Site VPN connections.

B. Deploy a transit gateway and a new accelerated Site-to-Site VPN connection.

C. Replace the existing Site-to-Site VPN connections with new Site-to-Site VPN connections that use IPv6.

D. Replace the existing Site-to-Site VPN connections with AWS PrivateLink connections.

Answer: B

Explanation:

Transit Gateway for Centralized Routing: A transit gateway centralizes the management and routing of multiple Site-to-Site VPN connections. It helps optimize network paths and reduce latency by avoiding the need for direct peering between each VPC and on-premises.

Accelerated Site-to-Site VPN: AWS accelerated Site-to-Site VPN connections use AWS Global Accelerator to reduce latency by leveraging the AWS global network to route traffic more efficiently. This significantly improves the performance of the applications.

Supports Existing IPv4 Configuration: Accelerated VPNs can work with IPv4 addresses, allowing the company to address latency without requiring migration to IPv6, which would introduce additional complexity.

Question: 274

A company has a transit gateway in a single AWS account. The company sends flow logs for the transit gateway to an Amazon CloudWatch Logs log group.

The company created an AWS Lambda function to analyze the logs. The Lambda function sends a notification to an Amazon Simple Notification Service (Amazon SNS) topic when a VPC generates traffic that is dropped by the transit gateway. Each notification contains the account ID. VPC ID, and total amount of dropped packets.

The company wants to subscribe a new Lambda function to the SNS topic. The new Lambda function must automatically prevent the traffic that is identified in each notification from leaving a VPC by applying a network ACL to the transit gateway attachment subnets in the VPC that generates the traffic.

Which solution will meet these requirements?

A. Configure the existing Lambda function to add the destination IP addresses of the dropped traffic to each SNS notification. Configure the new Lambda function to create an outbound rule by using the destination IP addresses in the network ACL.

B. Configure the existing Lambda function to add the source IP addresses of the dropped traffic to each SNS notification. Configure the new Lambda function to create an inbound rule by using the source IP addresses in the network ACL.

C. Configure the existing Lambda function to add the source IP addresses of the dropped traffic to each SNS notification. Configure the new Lambda function to create an outbound rule by using the source IP addresses in the network ACL.

D. Configure the existing Lambda function to add the destination IP addresses of the dropped traffic to each SNS notification. Configure the new Lambda function to create an inbound rule by using the destination IP addresses in the network ACL.

Answer: B

Explanation:

Source IP Addresses in Dropped Traffic: When traffic is dropped by the transit gateway, the source IP addresses are the origin of the traffic causing the issue. To block this traffic from entering the VPC, an inbound rule must be added to the network ACL for the transit gateway attachment subnets.

Inbound Rule on Network ACL: Network ACLs (NACLs) are stateless and require explicit rules to allow or deny traffic. Adding an inbound rule to deny traffic from the source IP addresses effectively prevents the unwanted traffic from entering the VPC.

Notification Details: The existing Lambda function should include the source IP addresses in the SNS notification so that the new Lambda function can use this information to automatically update the NACL.

Question: 275

A company has multiple VPCs with subnets that use IPv4. Traffic from the VPCs to the internet uses a NAT gateway. The company wants to transition to IPv6.

A network engineer creates multiple IPv6-only subnets in an existing testing VPC. The network engineer deploys a new Amazon EC2 instance that has an IPv6 address into one of the subnets. During testing, the network engineer discovers that the new EC2 instance is not able to communicate with an IPv4-only service through the internet. The network engineer needs to enable the IPv6 EC2 instance to communicate with the IPv4-only service.

Which solution will meet this requirement?

A. Enable DNS64 for the IPv6-only subnets. Update the route tables for the IPv6-only subnets to send traffic through the NAT gateway.

B. Enable NAT64 for the testing VPC. Reconfigure the existing NAT gateway to support IPv6.

C. Enable DNS64 for the new EC2 instance. Create a new egress-only internet gateway that supports IPv6.

D. Enable NAT64 for each route table. Create a new NAT gateway that supports both IPv4 and IPv6.

Answer: A

Explanation:

Understanding the Issue: The IPv6-only EC2 instance cannot communicate with IPv4-only services because IPv6 and IPv4 are not directly compatible. To bridge this gap, DNS64 and NAT64 are used together. However, AWS NAT gateways do not natively support NAT64, but you can use DNS64 to translate IPv4 DNS records (A records) into IPv6-compatible addresses (AAAA records).

DNS64 for IPv6-Only Subnets: DNS64 is a service that synthesizes AAAA records for IPv4-only services. This allows IPv6-only clients to resolve IPv4 addresses as IPv6-compatible addresses, enabling communication through the NAT gateway.

NAT Gateway with Route Table Updates: The NAT gateway enables outbound communication from private subnets to the internet. Updating the route tables for IPv6-only subnets to send traffic through the NAT gateway ensures that the IPv6 EC2 instance can reach IPv4 services.

Question: 276

A company deployed an application in two AWS Regions in one AWS account. The company has one VPC in each Region. The VPCs use non-overlapping private CIDR ranges.

The company needs to connect both VPCs to a single on-premises data center to test the application. The application requires up to 800 Mbps of throughput. A network engineer needs to establish connectivity between the VPCs and the on-premises data center.

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

A. Order a 2 Gbps Direct Connect connection for the data center. Configure a virtual private gateway in each VPC. Create a private VIF for each virtual private gateway, and associate the virtual private gateways with the Direct Connect connection. Configure static routes in the VPC route tables and in the data center router.

B. Order a 2 Gbps Direct Connect connection for the data center. Configure a virtual private gateway in each VPC. Create a private VIF for each virtual private gateway, and associate the virtual private gateways with the Direct Connect connection. Configure Open Shortest Path First (OSPF) routing between the private VIF and the data center.

C. Configure a customer gateway and a virtual private gateway in each VPC. Configure an AWS Site- to-Site VPN connection between the data center and each VPC. Configure static routes in each VPC route table to point to the subnets in the data center.

D. Configure a customer gateway and a virtual private gateway in each VPC. Configure an AWS Site- to-Site VPN connection between the data center and each VPC. Configure BGP routing between the VPCs and the data center.

Answer: A

Explanation:

Direct Connect for High Throughput: The application requires up to 800 Mbps of throughput. AWS Direct Connect is the best choice for meeting this requirement because it provides a high-bandwidth, low-latency, and private connection between the on-premises data center and AWS.

Private VIFs for VPC Connectivity: Private virtual interfaces (VIFs) are used to connect to VPCs through virtual private gateways. Configuring a private VIF for each VPC ensures secure and efficient connectivity to the on-premises data center.

Static Routes for Simplicity: Using static routes in the VPC route tables and the data center router minimizes operational overhead compared to dynamic routing protocols like BGP. Static routes are sufficient for testing purposes and are simpler to configure and maintain.

Question: 277

A company runs a workload in a single VPC on AWS. The company's architecture contains several interface VPC endpoints for AWS services, including Amazon CloudWatch Logs and AWS Key Management Service (AWS KMS). The endpoints are configured to use a shared security group. The security group is not used for any other workloads or resources.

After a security review of the environment, the company determined that the shared security group is more permissive than necessary. The company wants to make the rules associated with the security group more restrictive. The changes to the security group rules must not prevent the resources in the VPC from using AWS services through interface VPC endpoints. The changes must prevent unnecessary access.

The security group currently uses the following rules:

• Inbound - Rule 1

Protocol: TCP

Port: 443

Source: 0.0.0.0/0

• Inbound - Rule 2

Protocol: TCP

Port: 443

Source: VPC CIDR

• Outbound - Rule 1

Protocol: All

Port: All

Destination: 0.0.0.0/0

Which rule or rules should the company remove to meet with these requirements?

A. Outbound - Rule 2

B. Inbound - Rule 1 and Outbound - Rule 1

C. Inbound - Rule 2 and Outbound - Rule 1

D. Outbound - Rule 1

Answer: B

Explanation:

Inbound Rule 1 (Allow TCP 443 from 0.0.0.0/0): This rule allows all sources, including the public internet, to access the interface VPC endpoints. Since interface VPC endpoints are used within the VPC for communication with AWS services, this rule is unnecessarily permissive. Removing this rule enhances security while still allowing communication within the VPC using Rule 2 (TCP 443 from the VPC CIDR).

Outbound Rule 1 (Allow All Protocols, All Ports to 0.0.0.0/0): This rule is overly permissive and unnecessary for interface VPC endpoints, as traffic destined for AWS services through these endpoints does not need unrestricted outbound access. Removing this rule ensures that outbound traffic is limited to what is required for communication with the AWS services through the interface endpoints.

Question: 278

A company uses transit gateways to route traffic between the company's VPCs. Each transit gateway has a single route table. Each route table contains attachments and routes for the VPCs that are in the same AWS Region as the transit gateway. The route tables in each VPC also contain routes to all the other VPC CIDR ranges that are available through the transit gateways. Some VPCs route to local NAT gateways.

The company plans to add many new VPCs soon. A network engineer needs a solution to add new VPC CIDR ranges to the route tables in each VPC.

Which solution will meet these requirements in the MOST operationally efficient way?

A. Create a new customer-managed prefix list. Add all VPC CIDR ranges to the new prefix list. Update the route tables in each VPC to use the new prefix list ID as the destination and the appropriate transit gateway ID as the target.

B. Turn on default route table propagation for the transit gateway route tables. Turn on route propagation for each route table in each VPC.

C. Update the route tables in each VPC to use 0.0.0.010 as the destination and the appropriate transit gateway ID as the target.

D. Turn on default route table association for the transit gateway route tables. Turn on route propagation for each route

table in each VPC.

Answer: A

Explanation:

Using a Prefix List for Route Management: A customer-managed prefix list allows you to group multiple CIDR ranges into a single logical entity. By referencing the prefix list in VPC route tables, you can simplify route management. This eliminates the need to manually add individual CIDR ranges to each VPC route table.

Operational Efficiency: When a new VPC is added, its CIDR range can be added to the prefix list, and all route tables referencing the prefix list will automatically include the new CIDR. This reduces operational overhead compared to manually updating each route table.

Flexibility: The prefix list approach is highly scalable and supports the company's need to add many new VPCs in the future.

Question: 279

A company has several AWS Site-to-Site VPN connections between an on-premises customer gateway and a transit gateway. The company's application uses IPv4 to communicate through the VPN connections.

The company has updated the VPC to be dual stack and wants to transition to using IPv6-only for new workloads. When the company tries to communicate through the existing VPN connections, IPv6 traffic fails.

Which solution will provide IPv6 support with the LEAST operational overhead?

A. Create a new Site-to-Site VPN connection that supports IPv6.

B. Create a new Site-to-Site VPN connection to a self-managed Amazon EC2 instance that runs open source software.

C. Update the existing Site-to-Site VPN connections to support IPv6.

D. Update the on-premises customer gateway's public IP address from IPv4 to IPv6.

Answer: A

Explanation:

IPv6 Support in VPN Connections: Existing AWS Site-to-Site VPN connections that were originally configured for IPv4 do not automatically support IPv6 traffic. To enable IPv6 communication, a new Site-to-Site VPN connection must be created that explicitly supports IPv6.

Least Operational Overhead: Creating a new IPv6-enabled Site-to-Site VPN connection is straightforward and does not require extensive reconfiguration of the existing IPv4 setup. This ensures a smooth transition to dual-stack or IPv6-only workloads with minimal disruption.

Support for Dual-Stack Workloads: The new IPv6-enabled Site-to-Site VPN connection can coexist with the existing IPv4 connections, allowing the company to transition workloads incrementally to IPv6.

Question: 280

A company has two teams: Team A and Team B. Team A has VPCs that run in Account A. The team uses a transit gateway (TGW-A) to route traffic between workloads that run in the different VPCs.
Similarly, Team B has VPCs that run in Account B. Team B uses a different transit gateway (TGW-B) to route traffic between workloads that run in the different VPCs.
The company's network team manages the routing for Team A and Team B. The network team wants to retire TGW-B and use a single transit gateway to manage routing for the VPCs of both teams.

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

A. Create a resource share for TGW-A Share TGW-A with Account B. Create VPC attachments for the VPCs in Account B. Configure routing for the VPCs in TGW-A route tables. Update the route tables of the VPCs in Account B to forward traffic to TGW-A. Delete TGW-B attachments and TGW-B.

B. Create a resource share for TGW-A. Share TGW-A with Account B. Replicate the TGW-B configuration to TGW-A to automatically start routing changes for the VPCs in Account B. Delete TGW-B when routing changes are complete.

C. Create a new transit gateway (TGW-C) in Account A. Create a resource share for TGW-C. Share TGW-C with Account B. Create VPC attachments for the VPCs in Account A and Account B. Configure routing for all the VPCs in TGW-C route tables. Update the route tables for the VPCs in Account A and Account B to forward traffic to TGW-C. Delete TGW-A attachments and TGW-B attachments. Delete TGW-A and TGW-B.

D. Create a new transit gateway (TGW-C) in a new account (Account C). Create a resource share for TGW-C. Share TGW-C with Account A and Account B. Create VPC attachments for the VPCs in

Account A and Account B. Configure routing for all the VPCs in TGW-C route tables. Update the route tables for the VPCs in Account A and Account B to forward traffic to TGW-C. Delete TGW-A attachments and TGW-B attachments. Delete TGW-A and TGW-B.

Answer: A

Explanation:

This solution minimizes operational overhead by using a single transit gateway (TGW-A) for both teams, while also leveraging resource sharing between accounts. This approach eliminates the need to create new transit gateways, thus reducing complexity and the operational overhead of managing multiple transit gateways.

Question: 281

A company has an AWS environment that includes multiple VPCs that are connected by a transit gateway. The company wants to use a certificate-based AWS Site-to-Site VPN connection to establish connectivity between an on-premises environment and the AWS environment. The company does not have a static public IP address for the on-premises environment.

Which combination of steps should the company take to establish VPN connectivity between the transit gateway and the on-premises environment? (Choose two.)

A. Create a public certificate in AWS Certificate Manager (ACM).

B. Create a private certificate in AWS Certificate Manager (ACM).

C. Configure the Site-to-Site VPN tunnels to use the pre-shared key (PSK).

D. Create a customer gateway. Specify the current dynamic IP address of the customer gateway device's external interface.

E. Create a customer gateway. Do not specify the IP address of the customer gateway device.

Answer: B, D

Explanation:

Create a private certificate in AWS Certificate Manager (ACM): This involves setting up a private Certificate Authority (CA) within AWS ACM, which will be used to issue certificates for authenticating your customer gateway device.

Create a customer gateway. Specify the current dynamic IP address of the customer gateway device's external interface: Even though on-premises environment doesn't have a static IP, you can still configure the customer gateway in AWS by specifying its current dynamic IP address. This setup allows AWS to recognize and authenticate your customer gateway device during the VPN connection establishment.

Question: 282

A company operates in multiple AWS Regions. The company has deployed transit gateways in each Region. The company uses AWS Organizations to operate multiple AWS accounts in one organization.

The company needs to capture all VPC flow log data when a new VPC is created. The company needs to send flow logs to a

specific Amazon S3 bucket.

Which solution will meet these requirements with the LEAST administrative effort?

A. Update IAM permissions for each user to include a condition that ensures users can create VPCs only when VPC Flow Logs is enabled and configured correctly.

B. Create a custom AWS Config rule with automatic remediation that verifies VPC Flow Logs is enabled and configured correctly. Apply the AWS Config rule to the organization.

C. Enable VPC Flow Logs on each transit gateway. Configure VPC Flow Logs to send flow logs to the specified S3 bucket.

D. Deploy a serverless application that uses AWS CloudTrail to monitor for VPC creation events in each account. Configure the application to apply the correct VPC Flow Logs configuration.

Answer: B

Explanation:

This solution uses AWS Config, which allows you to automatically monitor and evaluate the configuration of AWS resources, including VPCs. By creating a custom AWS Config rule that checks whether VPC Flow Logs are enabled and correctly configured, you can ensure that VPC flow logs are captured for every new VPC. With automatic remediation, the rule can also ensure the VPC Flow Logs configuration is applied if not already set. Additionally, applying this rule to your entire organization will simplify the management process and reduce administrative effort.

Question: 283

A company wants to analyze TCP internet traffic. The traffic originates from Amazon EC2 instances in the company's VPC. The EC2 instances initiate connections through a NAT gateway.

The company wants to capture data about the traffic including source and destination IP addresses ports, and the first 8 bytes of the TCP segments of the traffic. The company needs to collect, store, and analyze all the required data points.

Which solution will meet these requirements?

A. Configure the EC2 instances to be VPC traffic mirror sources. Deploy software on the traffic mirror target to forward the data to Amazon CloudWatch Logs. Analyze the data by using CloudWatch Logs Insights

B. Configure the NAT gateway to be a VPC traffic mirror source. Deploy software on the traffic mirror target to forward the data to an Amazon S3 bucket. Analyze the data by using Amazon Athena.

C. Turn on VPC Flow Logs for the EC2 instances. Specify the default format and set Amazon CloudWatch Logs as the log destination. Analyze the flow log data by using CloudWatch Logs Insights.

D. Turn on VPC Flow Logs for the EC2 instances. Specify a custom format and set Amazon S3 as the log destination. Analyze

the flow log data by using Amazon Athena.

Answer: A

Explanation:

This solution meets the requirements for capturing detailed TCP internet traffic, including source and destination IP addresses, ports, and the first 8 bytes of TCP segments. By configuring the EC2 instances as traffic mirror sources and deploying a software solution on the target to forward the captured traffic to CloudWatch Logs, you can analyze the traffic in-depth using CloudWatch Logs Insights. VPC traffic mirroring is ideal for capturing low-level network traffic, providing the necessary data points for analysis.

Question: 284

A media company is planning to host an event that the company will live stream to users. The company wants to use Amazon CloudFront.

A network engineer creates a primary origin and a secondary origin for CloudFront. The engineer needs to ensure that the primary origin can fail over to the secondary origin within 15 seconds if a disruption occurs.

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

A. Configure a Lambda@Edge function to check the health status of both origins every 10 seconds. Reroute incoming requests when the origin health status is unhealthy.

B. Create a Network Load Balancer (NLB) in front of both origins Configure the NLB as the origin in CloudFront.

C. Set the CloudFront origin connection timeout value to 5 seconds Set the origin connection attempts value to 2.

D. Configure a Lambda@Edge function to monitor incoming requests for an origin response. Reroute incoming requests if no response is received from the primary origin within 10 seconds.

Answer: B

Explanation:

The solution involves using an NLB to manage the failover between the primary and secondary origins. The NLB automatically handles health checks for both origins and will route traffic to the healthy origin, providing a seamless failover within the required 15 seconds. This approach requires minimal operational overhead, as the NLB handles the routing and health checking without the need for custom code or manual intervention.

Question: 285

AnyCompany deploys and manages networking resources in its AWS network account, named Account-A. AnyCompany acquires Example Corp, which has an application that runs behind an Application Load Balancer (ALB) in Example Corp's AWS account, named Account-B.

Example Corp needs to use AWS Global Accelerator to create an accelerator to publish the application to users. AnyCompany's networking team will manage the accelerator.

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

A. Create an accelerator in Account-B. Use a cross-account role from Account-A to grant the networking team access to manage the accelerator.

B. Deploy a Network Load Balancer (NLB) in Account-A to route traffic to the ALB in Account-B. Create an accelerator, and set the NLB as the endpoint in Account-A.

C. Create a cross-account Global Accelerator attachment in Account-B for the Account-A principal. Create an accelerator in Account-A by using the shared attachment.

D. Create an accelerator in Account-A. Use AWS Resource Access Management (AWS RAM) to share the accelerator with Account-B. Associate the ALB in Account-B with the accelerator in Account-A.

Answer: C

Explanation:

Chosen solution involves setting up a cross-account Global Accelerator attachment from Account-B (where the application is hosted) to Account-A (where the accelerator will be managed). By using this shared attachment, the networking team in Account-A can manage the Global Accelerator with minimal management overhead, while still allowing the ALB in Account-B to be the endpoint for the accelerator. This approach requires fewer resources and minimizes complexity compared to other solutions.

Question: 286

A company has two AWS Direct Connect connections between Direct Connect locations and the company's on-premises environment in the US. The company uses the connections to communicate with AWS workloads that run in the us-east-1 Region. The company has a transit gateway that connects several VPCs. The Direct Connect connections terminate at a Direct Connect gateway and the transit VIFs to the transit gateway.

The company recently acquired a smaller company that is based in Europe. The newly acquired company has only on-premises workloads. The newly acquired company does not

expect to run workloads on AWS for the next 3 years. However, the newly acquired company requires connectivity to the parent company's AWS resources in us-east-1 and to the

parent company's on-premises environment in the US. The parent company wants to use two new Direct Connect connections in Europe to provide the required connectivity.

Which solution will meet these requirements with the LEAST operational overhead for the newly acquired company?

A. Associate new transit VIFs to the existing Direct Connect gateway. Configure the new transit VIFs to use Direct Connect SiteLink.

B. Associate new transit VIFs to a new Direct Connect gateway and to a new transit gateway in the eu-west-1 Region. Use transit gateway peering to connect the transit gateways.

C. Associate new private VIFs to the existing Direct Connect gateway. Configure the existing transit VIFs and the new private VIFs to use Direct Connect SiteLink.

D. Associate new private VIFs to a new Direct Connect gateway and to a new VPC in us-east-1. Configure the existing transit VIFs and the new private VIFs to use Direct Connect SiteLink and AWS PrivateLink endpoints in the new VPC.

Answer: A

Explanation:

In this scenario, the company wants to provide connectivity from the newly acquired company in Europe to the existing AWS resources in the us-east-1 Region with minimal operational overhead.

The best approach is to use Direct Connect SiteLink, which allows direct communication between two different Direct Connect locations (one in Europe and one in the US) via the existing Direct Connect gateway.

By associating new transit VIFs (Virtual Interfaces) to the existing Direct Connect gateway and configuring Direct Connect SiteLink, the company can efficiently extend the existing network architecture with minimal additional configuration. This solution provides the required connectivity to both AWS resources and the on-premises environment in the US, leveraging the existing infrastructure without introducing significant complexity or the need for additional resources like new transit gateways or VPCs.

Question: 287

A company is establishing hybrid cloud connectivity from an on-premises environment to AWS in the us-east-1 Region. The company is using a 10 Gbps AWS Direct Connect dedicated connection. The company has two accounts in AWS. Account A has transit gateways in four AWS Regions. Account B has transit gateways in three Regions. The company does not plan to expand.

To meet security requirements the company's accounts must have separate cloud infrastructure.

Which solution will meet these requirements MOST cost-effectively?

A. Create one Direct Connect gateway in us-east-1. Use AWS Resource Access Manager (AWS RAM) to share the Direct Connect gateway with each account. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway. Create a transit VIF for Account

B. Associate the three transit gateways in Account B to the Direct Connect gateway.

B. Create one Direct Connect gateway in us-east-1 for Account A. Create a second Direct Connect gateway in us-east-1 for Account B. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway in Account A. Create a transit VIF for Account B. Associate the three transit gateways in Account B to the Direct Connect gateway in Account B.

C. Create one Direct Connect gateway in us-east-1. Use AWS Resource Access Manager (AWS RAM) to share the Direct Connect gateway with each account. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway. Order a new 10 Gbps Direct Connect dedicated connection for Account B. Create a transit VIF on the new Direct Connect connection for Account B. Associate the three transit gateways in Account B to the Direct Connect gateway.

D. Create one Direct Connect gateway in us-east-1 for Account A. Create a second Direct Connect gateway in us-east-1 for Account B. Create a transit VIF for Account A. Associate the four transit gateways in Account A to the Direct Connect gateway in Account A. Order a new 10 Gbps Direct Connect dedicated connection for Account B. Create a transit VIF on the new Direct Connect connection for Account B. Associate the three transit gateways in Account B to the Direct Connect gateway in Account B.

Answer: A

Explanation:

The most cost-effective and scalable solution is to create a single Direct Connect gateway in us-east- 1, and use AWS Resource Access Manager (AWS RAM) to share the Direct Connect gateway between Account A and Account B. This approach avoids the need for multiple Direct Connect connections and allows both accounts to share the same connection, which is a more cost-efficient solution compared to creating separate connections for each account.

Transit VIFs (Virtual Interfaces) will be created for both Account A and Account B, and each account's respective transit gateways will be associated with the same Direct Connect gateway. This solution allows both accounts to access AWS resources in the most efficient manner.

Question: 288

A company runs an application across multiple AWS Regions and multiple Availability Zones. The company needs to expand to a new AWS Region. Low latency is critical to the functionality of the application.

A network engineer needs to gather metrics for the latency between the existing. Regions and the new Region. The network engineer must gather metrics for at least the previous 30 days.

Which solution will meet these requirements?

A. Configure an AWS Network Access Analyzer Network Access Scope, and use the analysis to review the latency.

B. Set up AWS Network Manager Infrastructure Performance. Publish network performance metrics to Amazon CloudWatch.

C. Use an Amazon VPC Reachability Analyzer path to review the latency.

D. Set up VPC Flow Logs. Publish log metrics to Amazon CloudWatch.

Answer: B

Explanation:

AWS Network Manager Infrastructure Performance is specifically designed to monitor the network performance across multiple AWS Regions, and it provides network metrics, including latency, between AWS Regions. By setting it up and publishing the metrics to Amazon CloudWatch, the network engineer can gather the necessary latency metrics for at least the previous 30 days. This solution directly addresses the requirement for low latency and monitoring network performance between the existing Regions and the new Region.

Question: 289

A company operates in the us-east-1 Region and the us-west-1 Region. The company is designing a solution to connect an on-premises data center to the company's AWS environment in us-east-1. The solution uses two AWS Direct Connect connections.

Traffic from us-west-1 to the data center needs to traverse the Direct Connect connections. A network engineer needs to set up active-passive functionality across the two Direct Connect connections by using a Direct Connect gateway to influence inbound traffic from VPCs that are in us- west-1 to the data center.

Which solution will meet these requirements?

A. At the data center, set the local preference for the primary connection to be higher than the local preference for the secondary connection.

B. Use AS path prepending to set the AS path on the primary connection to be longer than the AS path on the secondary connection.

C. Use local preference BGP community tags to apply the 7224:7300 local preference BGP community tag to the prefixes for the primary connection. Apply the 7224:7100 local preference BGP community tag to the prefixes for the secondary connection.

D. Use local preference BGP community tags to apply the 7224:9300 local preference BGP community tag to the prefixes for the primary connection. Apply the 7224:9100 local preference BGP community tag to the prefixes for secondary connection.

Answer: D

Explanation:

To control inbound traffic from AWS to the on-premises data center, local preference BGP community tags are used with a Direct Connect gateway.

AWS uses the following tags:

-7224:9300 sets higher local preference (preferred route).

■ 7224:9100 sets lower local preference (less preferred).

Applying 7224:9300 to the primary connection and 7224:9100 to the secondary connection ensures active-passive routing behavior for inbound traffic from AWS VPCs.

Question: 290

A company has multiple firewalls and ISPs for its on-premises data center. The company has a single AWS Site-to-Site VPN connection from the company's on-premises data center to a transit gateway. A single ISP services the Site-to-Site VPN connection. Multiple VPCs are attached to the transit gateway.

A customer gateway that the Site-to-Site VPN connection uses fails. Connectivity is completely lost, but the company's network team does not receive a notification.

The network team needs to implement redundancy within a week in case a single customer gateway fails again. The team wants to use an Amazon CloudWatch alarm to send notifications to an Amazon Simple Notification Service (Amazon SNS) topic if any tunnel of the Site-to-Site VPN connection fails.

Which solution will meet these requirements MOST cost-effectively?

A. Replace the existing customer gateway with a new router. Create a new Site-to-Site VPN connection to the transit gateway. For each VPN connection, set up a CloudWatch TunnelState alarm for the VPN connection. Use a value of 0 for the alarm.

B. Use a second customer gateway and a second ISP. Create a new Site-to-Site VPN connection to the transit gateway. For each VPN connection, set up a CloudWatch TunnelState alarm for the VPN connection. Use a value of less than 1 for the alarm.

C. Add an AWS Direct Connect connection to the existing Site-to-Site VPN connection to the transit gateway. For each VPN connection, set up a CloudWatch TunnelState alarm for the VPN connection. Use a value of failed for the alarm.

D. Use a second customer gateway with the existing ISP. Create a new Site-to-Site VPN connection to the transit gateway. For each VPN connection, set up a CloudWatch TunnelState alarm for the VPN connection. Use a value of unavailable for the alarm.

Answer: B

Explanation:

Redundancy requires a second customer gateway and ideally a second ISP to avoid a single point of failure. AWS Site-to-Site VPN connections support two tunnels per VPN connection.

By creating a second VPN connection (to the transit gateway) with a second customer gateway and ISP, the solution meets the redundancy requirement.

CloudWatch TunnelState alarms can be configured on each tunnel. A value of < 1 (i.e., when the tunnel is down) will trigger the alarm.

Question: 291

A company ran out of IP address space in one of the Availability Zones in an AWS Region that the company uses. The Availability Zone that is out of space is assigned the

10.10.1.0/24 CIDR block. The company manages its networking configurations in an AWS CloudFormation stack. The company's VPC is assigned the 10.10.0.0/16 CIDR

block and has available capacity in the 10.10.1.0/22 CIDR block.

How should a network specialist add more IP address space in the existing VPC with the LEAST operational overhead?

A. Update the AWS :: EC2 :: Subnet resource for the Availability Zone in the CloudFormation stack. Change the CidrBlock property to 10.10.1.0/22.

B. Update the AWS :: EC2 :: VPC resource in the CloudFormation stack. Change the CidrBlock property to 10.10.1.0/22.

C. Copy the CloudFormation stack. Set the AWS :: EC2 :: VPC resource CidrBlock property to 10.10.0.0/16. Set the AWS :: EC2 :: Subnet resource CidrBlock property to 10.10.1.0/22 for the Availability Zone.

D. Create a new AWS :: EC2 :: Subnet resource for the Availability Zone in the CloudFormation stack. Set the CidrBlock property to 10.10.2.0/24.

Answer: D