Free Questions for ANS-C01 by certsdeals

Shared by Bradshaw on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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?

Options:

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 2

Question Type: MultipleChoice

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?

Options:

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 3

Question Type: **MultipleChoice**

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?

Options:

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 4

Question Type: MultipleChoice

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?

Options:

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 5

Question Type: **MultipleChoice**

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?

Options:

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 6

Question Type: MultipleChoice

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?

Options:

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 7

Question Type: MultipleChoice

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.)

Options:

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:

B, C

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 PrivateLink 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 8

Question Type: MultipleChoice

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.)

Options:

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:

B, C, F

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 cross-Region 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 9

Question Type: MultipleChoice

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.)

Options:

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:

D

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.

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.

Question 10

Question Type: **MultipleChoice**

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?

Options:

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 11

Question Type: MultipleChoice

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?

Options:

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

Question 12

Question Type: MultipleChoice

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.)

Options:

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:

B, D, E