Answer Sheet

Q1) Your company currently has a VPC defined as 10.0.0.0/16. Subnets are defined in this VPC along with Instances created in the subnet. You need to ensure that resources in the VPC can resolve your on-premise DNS resources. How can you achieve this? Choose 2 answers from the options given below.

● Configure DHCP Options for your Subnet to point to the EC2 Instance.
● Create a private hosted zone in Route53
✓ Configure DHCP Options for your VPC to point to the EC2 Instance.
Explanation:-Here you can create your own EC2 Instance which will act as the DNS server. The VPC can then use the DHCP Options which points to this EC2 Instance as the DNS resolver. For more information on DNS in a VPC , please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-dns.html
✓ Create an EC2 Instance in your VPC which will act as the DNS server
Explanation:-Here you can create your own EC2 Instance which will act as the DNS server. The VPC can then use the DHCP Options which points to this EC2 Instance as the DNS resolver. For more information on DNS in a VPC , please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-dns.html

---

Q2) You need to setup EC2 instances inside a VPC. The requirement is also to create a standby interface in case of any of the EC2 instances not responding to traffic. How can you achieve this?

✓ Attach a secondary ENI to the Instance
Explanation:-AWS Docs provides following details: Scenarios for Network Interfaces Attaching multiple network interfaces to an instance is useful when you want to: Create a management network. Use network and security appliances in your VPC. Create dual-homed instances with workloads/roles on distinct subnets. Create a low-budget, high-availability solution. Creating a Management Network You can create a management network using network interfaces. In this scenario, the primary network interface (eth0) on the instance handles public traffic and the secondary network interface (eth1) handles backend management traffic and is connected to a separate subnet in your VPC that has more restrictive access controls. The public interface, which may or may not be behind a load balancer, has an associated security group that allows access to the server from the internet (for example, allow TCP port 80 and 443 from 0.0.0.0/0, or from the load balancer) while the private facing interface has an associated security group allowing SSH access only from an allowed range of IP addresses either within the VPC or from the internet, a private subnet within the VPC or a virtual private gateway. To ensure failover capabilities, consider using a secondary private IPv4 for incoming traffic on a network interface. In the event of an instance failure, you can move the interface and/or secondary private IPv4 address to a standby instance. For more information on using the Elastic Network Interface, please refer to the below URL https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html
● Attach an elastic IP to the Instance
● Attach a public and private IP to the instance
● Assign a secondary IP to the ENI attached to the EC2 Instance

---

Q3) You're planning on creating a VPN connection to 2 VPC's in AWS. You are going to be using the same customer gateway in both cases. These VPC's have overlapping CIDR blocks. What can be done to ensure the routing is done right on the customer side.

● Use static routes on the customer side
● Configure AS_PATH for each of the routes
✓ Use VRF technology for routing
Explanation:-This is given in the AWS Documentation Virtual Routing and Forwarding (VRF) is a technology that allows multiple instances of a routing table to co-exist within the same router at the same time. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other. AWS recommends using VRFs when connecting a single customer gateway to multiple Amazon VPCs because the VPN connection creation logic is designed to ensure unique tunnel IP addresses for each connection within a single VPC, but not necessarily across multiple VPCs. For more information on configuring routes to multiple VPC's , please refer to the below URL https://aws.amazon.com/articles/connecting-a-single-customer-router-to-multiple-vpcs/
● Use BFD technology for routing

Q4) You have created a NAT gateway to ensure that instances in your private subnet can download updates from the internet. But the instances are still not able to reach the internet even after the gateway has been created. Which of the following could be one of the underlying issue?

✓ The NAT gateway has been created in the private subnet
Explanation:-The AWS Documentation mentions the following To troubleshoot instances that can't connect to the Internet from a private subnet using a NAT gateway, check the following: Verify that the destination is reachable by pinging the destination from another source using a public IP address. Verify that the NAT gateway is in the Available state. Note: A NAT gateway in the Failed state is automatically deleted after about an hour. Make sure that you've created your NAT gateway in a public subnet, and that that the public route table has a default route pointing to an Internet gateway. Make sure that the private subnet's route table has a default route pointing to the NAT gateway. Check that you have allowed the required protocols and ports for outbound traffic to the Internet. For more information on NAT gateways , please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-nat-gateway.html
● The NAT gateway has not been created with the wrong AMI
● The NAT gateway has been created in the public subnet
● The NAT gateway has been created with the wrong Instance type

---

Q5) You have an application that consists of the following setup • An EC2 Instance that supports the main front end part of the application • An EC2 Instance that is used to process Images You are planning on using a load balancer to route requests based on the type of request and then route them to the respective servers. How can you accomplish this? Choose 2 answers from the options given below

● Create a Classic load balancer
✓ Create an Application load balancer
Explanation:-Here you need to route traffic based on the type of URL request. So based on the URL request , the request could go to either EC2 Instance. For this you need to create an Application Load balancer and target groups for each EC2 Instance For more information on Application

Load Balancers , please refer to the below URL https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html

✓ Create different target groups

Explanation:-Here you need to route traffic based on the type of URL request. So based on the URL request , the request could go to either EC2 Instance. For this you need to create an Application Load balancer and target groups for each EC2 Instance For more information on Application Load Balancers , please refer to the below URL https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html

● Create a TCP listener

Q6) Your company is planning on setting up applications over the internet with the following aspects • 2 Applications , each with their own domain name • Each application will have EC2 Instances as Web servers You need to ensure High Availability for the servers and also configure Route53. How would you achieve this? Choose 2 answers from the options given below

● Create a private Elastic Load Balancer

● Configure 2 private hosted zones in Route 53

✓ Configure 2 public hosted zones in Route 53

Explanation:-Since 2 domains are required , and hence since these are Web servers which most probably will need to be exposed to the Internet, you need to define 2 separate public hosted zones and an ELB. For more information on working with Hosted Zones , please refer to the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHZWorkingWith.html

✓ Create a public Elastic Load Balancer

Explanation:-Since 2 domains are required , and hence since these are Web servers which most probably will need to be exposed to the Internet, you need to define 2 separate public hosted zones and an ELB. For more information on working with Hosted Zones , please refer to the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHZWorkingWith.html

Q7) Your company has set EC2 Instances in a VPC. These Instances have been configured to query an on-premise Data center DNS server. But the Instances are not able to reach the On-premise server. Which of the following could be a reason for this? Choose 2 answers from the options given below

✓ The NACL's are blocking outgoing on port 53 for UDP

Explanation:-In order to communicate with a DNS Server , the instance needs to reach the DNS server on port 53 for both TCP and UDP. For more information on NACL's , please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

✓ The NACL's are blocking outgoing on port 53 for TCP

Explanation:-In order to communicate with a DNS Server , the instance needs to reach the DNS server on port 53 for both TCP and UDP. For more information on NACL's , please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

● The NACL's are blocking incoming on port 53 for TCP

● The NACL's are blocking incoming on port 53 for UDP

Q8) Your company is planning on setting up an application that consists of EC2 Instances , an Application Load Balancer and Cloudfront. Your management is worried about DDOs attacks. Which of the following can help protect against such network attacks? Choose 3 answers from the options given below

✓ Place the AWS WAF in front of the Application Load Balancer

Explanation:-The AWS Documentation mentions the following AWS WAF is a web application firewall that lets you monitor the HTTP and HTTPS requests that are forwarded to Amazon CloudFront or an Application Load Balancer. AWS WAF also lets you control access to your content. Based on conditions that you specify, such as the IP addresses that requests originate from or the values of query strings, CloudFront or an Application Load Balancer responds to requests either with the requested content or with an HTTP 403 status code (Forbidden). AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, and Route 53 hosted zones. AWS Shield Advanced incurs additional charges. For more information on AWS WAF , please refer to the below URL https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html

✓ Place the AWS WAF in front of the Cloudfront Distribution

Explanation:-The AWS Documentation mentions the following AWS WAF is a web application firewall that lets you monitor the HTTP and HTTPS requests that are forwarded to Amazon CloudFront or an Application Load Balancer. AWS WAF also lets you control access to your content. Based on conditions that you specify, such as the IP addresses that requests originate from or the values of query strings, CloudFront or an Application Load Balancer responds to requests either with the requested content or with an HTTP 403 status code (Forbidden). AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, and Route 53 hosted zones. AWS Shield Advanced incurs additional charges. For more information on AWS WAF , please refer to the below URL https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html

● Place the AWS WAF in front of the EC2 Instances

✓ Consider using AWS Shield Advanced

Explanation:-The AWS Documentation mentions the following AWS WAF is a web application firewall that lets you monitor the HTTP and HTTPS requests that are forwarded to Amazon CloudFront or an Application Load Balancer. AWS WAF also lets you control access to your content. Based on conditions that you specify, such as the IP addresses that requests originate from or the values of query strings, CloudFront or an Application Load Balancer responds to requests either with the requested content or with an HTTP 403 status code (Forbidden). AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, CloudFront distributions, and Route 53 hosted zones. AWS Shield Advanced incurs additional charges. For more information on AWS WAF , please refer to the below URL https://docs.aws.amazon.com/waf/latest/developerguide/what-is-aws-waf.html

Q9) Your company is planning on using Cloudfront along with S3 as the origin. There is a requirement to serve private content from S3. There is a requirement to ensure that access is restricted for certain individual files. How would you deliver the private content.

✓ Use Signed URL's

Explanation:-The AWS Documentation mentions the following Use signed URLs in the following cases: • You want to use an RTMP distribution. Signed cookies aren't supported for RTMP distributions. • You want to restrict access to individual files, for example, an installation download for your application. • Your users are using a client (for example, a custom HTTP client) that doesn't support cookies. Use signed cookies in the following cases: • You want to provide access to multiple restricted files, for example, all of the files for a video in HLS format or all of the files in the subscribers' area of a website. • You don't want to change your current URLs. For more information on a better understanding on serving private content , please refer to the below URL https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-choosing-signed-urls-cookies.html

● Use Private Keys

● Use Security Groups
● Use Signed Cookies

Q10) You have a set of EC2 Instances that are deployed in a VPC. An application is hosted on these instances. There are some issues which keep on recurring in the application and you plan to inspect the packets being sent from the application to trace the error. How can you achieve this?

● Use VPC Flow logs
✓ Use an IDS
Explanation:-Here you will need a custom Intrusion Detection system to do a packet level analysis. For more information on IDS, please refer to the below URL https://aws.amazon.com/mp/scenarios/security/ids/
● Use Cloudtrail
● Use Cloudwatch Logs

Q11) You have a requirement to ensure that hosted zones created in Route 53 have name servers that resonate with your domain name. How can you achieve this? Choose 2 answers from the options given below

✓ Create a Reusable delegation set using the AWS CLI
Explanation:-The AWS Documentation mentions the following Each Amazon Route 53 hosted zone is associated with four name servers, known collectively as a delegation set. By default, the name servers have names like ns-2048.awsdns-64.com. If you want the domain name of your name servers to be the same as the domain name of your hosted zone, for example, ns1.example.com, you can configure white label name servers, also known as vanity name servers or private name servers. To create a reusable delegation set, you can use the Route 53 API, the AWS CLI, or one of the AWS SDKs. For more information on reusable delegation set, please refer to the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/white-label-name-servers.html
● Create a Reusable delegation set using the AWS Console
✓ Create a Reusable delegation set using Route 53 API's
Explanation:-The AWS Documentation mentions the following Each Amazon Route 53 hosted zone is associated with four name servers, known collectively as a delegation set. By default, the name servers have names like ns-2048.awsdns-64.com. If you want the domain name of your name servers to be the same as the domain name of your hosted zone, for example, ns1.example.com, you can configure white label name servers, also known as vanity name servers or private name servers. To create a reusable delegation set, you can use the Route 53 API, the AWS CLI, or one of the AWS SDKs. For more information on reusable delegation set, please refer to the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/white-label-name-servers.html
● Specify the domain name when creating the record set for the name servers

Q12) Your account is part of a parent AWS Account that has a Direct Connect connection. You plan to use the AWS Direct Connection as a hosted VIF. What would you get charged for?

● The port hours
● The initial connection charges
● All Data transfer in
✓ All Data transfer out
Explanation:-For more information on AWS Direct connect pricing, please refer to the below URL https://aws.amazon.com/directconnect/pricing/

Q13) You have 2 VPC's VPCA(172.16.0.0/16) and VPCB(10.0.0.0/16). You are planning on establishing VPC connecting peering. Which of the following routes need to be added to the route table for both VPC's to ensure communication across VPC's. Choose 2 answers from the options given below. Assume that the Target for the VPC Peering connection has an ID of pcx-1122

✓ In the Route table for VPCB add a route of 172.16.0.0/16 and Target as pcx-1122
Explanation:-For more information on this example , one can visit the below URL http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html
● In the Route table for VPCB add a route of 10.0.0.0/16 and Target as pcx-1122
✓ In the Route table for VPCA add a route of 10.0.0.0/16 and Target as pcx-1122
Explanation:-For more information on this example , one can visit the below URL http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html
● In the Route table for VPCA add a route of 172.16.0.0/16 and Target as pcx-1122

Q14) You need to have instances created in a VPC which can support network speeds of upto 20 Gbps. Which of the following would be part of your implementation steps? Choose 2 answers from the options given below

● Create an Instance from an Instance type that supports the Intel 82599 VF interface
✓ Create an Instance from an Instance type that supports Enhanced Networking
Explanation:-For speeds for up to 25 Gbps , you need to choose an Instance type that supports Enhanced Networking Also ensure that Enhanced Networking is enabled on the device. For more information on Enhanced Networking , one can visit the below URL https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/enhanced-networking.html
✓ Enable Enhanced Networking if not already done
Explanation:-For speeds for up to 25 Gbps , you need to choose an Instance type that supports Enhanced Networking Also ensure that Enhanced Networking is enabled on the device. For more information on Enhanced Networking , one can visit the below URL https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/enhanced-networking.html
● Place the Instances in a placement group

Q15) You currently have set up a VPC, route tables with routes defined for traffic and Subnets in AWS. You just want to establish communication across all hosts. But you notice that some applications are not working as desired. These are Ipv6 based applications that are sitting across subnets in the VPC. What must be done to alleviate this issue?

● Ensure that the route of 0.0.0.0/0 is removed and a more specific route is placed.
● Remove the route of 0.0.0.0/0 and add the route of ::/0 instead to allow all communication.
✓ Add a route for ::/0 to the route table as well.

Explanation:-The AWS Documentation mentions the following CIDR blocks for IPv4 and IPv6 are treated separately. For example, a route with a destination CIDR of 0.0.0.0/0 (all IPv4 addresses) does not automatically include all IPv6 addresses. You must create a route with a destination CIDR of ::/0 for all IPv6 addresses. For more information on Route propagation , one can visit the below URL http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Route_Tables.html

● Add the default route of 172.132.0.0/16 to the Route table

Q16) When configuring Active Passive configuration for your VPN connections which of the following can be used to achieve such a configuration. Choose 2 answers from the options given below

✓ Use AS_PATH prepending

Explanation:-The AWS Documentation mentions this multi-data-center-config Redundant Active/Active VPN Connections Many AWS customers choose to implement VPN connections because they can be a quick, easy, and cost-effective way to set up remote connectivity to a VPC. To enable redundancy, each AWS Virtual Private Gateway (VGW) has two VPN endpoints with capabilities for static and dynamic routing. Although statically routed VPN connections from a single customer gateway are sufficient for establishing remote connectivity to a VPC, this is not a highly available configuration. The best practice for making VPN connections highly available is to use redundant customer gateways and dynamic routing for automatic failover between AWS and customer VPN endpoints. For simplicity, the diagram in the next section depicts each VPN connection, consisting of two IPsec tunnels to both VGW endpoints, as a single line. Configuration Details The configuration in this example consists of four fully meshed, dynamically routed IPsec tunnels between both VGW endpoints and two customer gateways. AWS provides configuration templates for a number of supported VPN devices to assist in establishing these IPsec tunnels and configuring BGP for dynamic routing. In addition to the AWS-provided VPN and BGP configuration details, customers must configure VPCs to efficiently route traffic to their data center networks. In this example, the VGW will prefer to send 10.0.0.0/16 traffic to Data Center 1 through Customer Gateway 1, and only reroute this traffic through Data Center 2 if the connection to Data Center 1 is down. Likewise, 10.1.0.0/16 traffic will prefer the VPN connection originating from Data Center 2. AWS recommends using one of the following approaches for communicating these route preferences (For a full explanation of VPC routing rule algorithm, see Configuring Multiple VPN Connections to Your Amazon VPC): More specific routes: With this approach, both Customer Gateway 1 and Customer Gateway 2 advertise a summary route of 10.0.0.0/15. In addition, Customer Gateway 1 advertises 10.0.0.0/16 and Customer Gateway 2 advertises 10.1.0.0/16. AWS will use the more specific routes to send traffic to the appropriate data center, and will fail back to the other data center following the summarized route if the more specific route becomes temporarily unavailable. AS-path prepending: With this approach, both Customer Gateway 1 and Customer Gateway 2 advertise 10.0.0.0/16 and 10.1.0.0/16. However, Customer Gateway 1 uses AS-path prepending when advertising the 10.1.0.0/16 network to make this route less preferred. Likewise, Customer Gateway 2 uses AS-path prepending when advertising the 10.0.0.0/16 network to make this route less preferred. AWS will use the preferred routes to send traffic to the appropriate data center, and will fail back to the other data center following the less preferred routes when necessary. For more information on High Availability Network connections , one can visit the below URL https://aws.amazon.com/answers/networking/aws-multiple-data-center-ha-network-connectivity/

✓ Use more specific routes

Explanation:-The AWS Documentation mentions this multi-data-center-config Redundant Active/Active VPN Connections Many AWS customers choose to implement VPN connections because they can be a quick, easy, and cost-effective way to set up remote connectivity to a VPC. To enable redundancy, each AWS Virtual Private Gateway (VGW) has two VPN endpoints with capabilities for static and dynamic routing. Although statically routed VPN connections from a single customer gateway are sufficient for establishing remote connectivity to a VPC, this is not a highly available configuration. The best practice for making VPN connections highly available is to use redundant customer gateways and dynamic routing for automatic failover between AWS and customer VPN endpoints. For simplicity, the diagram in the next section depicts each VPN connection, consisting of two IPsec tunnels to both VGW endpoints, as a single line. Configuration Details The configuration in this example consists of four fully meshed, dynamically routed IPsec tunnels between both VGW endpoints and two customer gateways. AWS provides configuration templates for a number of supported VPN devices to assist in establishing these IPsec tunnels and configuring BGP for dynamic routing. In addition to the AWS-provided VPN and BGP configuration details, customers must configure VPCs to efficiently route traffic to their data center networks. In this example, the VGW will prefer to send 10.0.0.0/16 traffic to Data Center 1 through Customer Gateway 1, and only reroute this traffic through Data Center 2 if the connection to Data Center 1 is down. Likewise, 10.1.0.0/16 traffic will prefer the VPN connection originating from Data Center 2. AWS recommends using one of the following approaches for communicating these route preferences (For a full explanation of VPC routing rule algorithm, see Configuring Multiple VPN Connections to Your Amazon VPC): More specific routes: With this approach, both Customer Gateway 1 and Customer Gateway 2 advertise a summary route of 10.0.0.0/15. In addition, Customer Gateway 1 advertises 10.0.0.0/16 and Customer Gateway 2 advertises 10.1.0.0/16. AWS will use the more specific routes to send traffic to the appropriate data center, and will fail back to the other data center following the summarized route if the more specific route becomes temporarily unavailable. AS-path prepending: With this approach, both Customer Gateway 1 and Customer Gateway 2 advertise 10.0.0.0/16 and 10.1.0.0/16. However, Customer Gateway 1 uses AS-path prepending when advertising the 10.1.0.0/16 network to make this route less preferred. Likewise, Customer Gateway 2 uses AS-path prepending when advertising the 10.0.0.0/16 network to make this route less preferred. AWS will use the preferred routes to send traffic to the appropriate data center, and will fail back to the other data center following the less preferred routes when necessary. For more information on High Availability Network connections , one can visit the below URL https://aws.amazon.com/answers/networking/aws-multiple-data-center-ha-network-connectivity/

● Use IPSec routing

● Use different ASN numbers

Q17) You have two Direct Connect connections and two VPN connections to your network. Following are the details Site A is VPN 10.2.0.0/24 7224:7100 AS 65000 65000 Site B is VPN 10.2.0.252/30 7224:7300 AS 65000 Site C is DX 10.0.0.0/8 AS 7224:7100 65000 65000 Site D is DX 10.0.0.0/16 AS 7224:7100 65000 65000 65000 Which site will AWS choose to reach your network?

● Site A

✓ Site B

Explanation:-7224:7300 denotes "High preference" and also has the longest prefix. Hence all other options by default become invalid. For more information on Route tables , one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Route_Tables.html Please refer to the below link on page 4 https://docs.aws.amazon.com/directconnect/latest/UserGuide/dc-ug.pdf

● Site D

● Site C

Q18) You've setup VPC Flow logs for your EC2 Instance ENI in a subnet. You can see the below REJECT record in the VPC Flow logs. What does this indicate. 2 123456789911 eni-abc123de 172.31.9.69 172.31.9.12 49761 3389 6 20 4249 1418530010 1418530070 REJECT OK

● A request was made on port 80 to the Instance

● Someone was trying to log into the Instance via SSH

✓ Someone was trying to log into the Instance via RDP

Explanation:-In the record which is recorded in VPC Flow logs , the highlighted field shown below shows that a request was made to port 3389 which is the RDP protocol 2 123456789911 eni-abc123de 172.31.9.69 172.31.9.12 49761 3389 6 20 4249 1418530010 1418530070 REJECT OK

There are also other possibilities such as a "port sniffer" that tried to use port 3389 (for RDP) or any other application that might have established contact with port 3389 (using cross-scripting for attacks). Here 'someone' might refer to a individual user (hacker) or even other malicious application to gain entry through the backdoor. By default all other options becomes invalid since clearly the log shows what is the port number recorded. For more information on VPC Flow Logs , one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html

● A request was made on port 443 to the Instance

Q19) Your company has the following setup with dynamic routing at customer gateway The routes advertised by both Customer gateway 1 and Customer Device 2 is 10.0.0.0/16. How will the traffic flow from the AWS to data center?

● The traffic will flow primarily through Customer gateway 1

✓ The traffic will flow primarily through Customer device 2

Explanation:-This sort of scenario is given in the AWS Documentation Configuration Details The configuration in this example consists of two dynamically routed connections, one using AWS Direct Connect and the other using a VPN connection from two different customer devices. AWS provides example router configurations to assist in establishing both AWS Direct Connect and VPN connections with BGP for dynamic routing. By default, AWS will always prefer to send traffic over an AWS Direct Connect connection, so no additional configuration is required to define primary and backup connections. In this example, both Customer Gateway 1 and Customer Device 2 advertise a summary route of 10.0.0.0/16 and AWS will send all traffic to Customer Device 2 as long as this network path is available. By default all other options become invalid For more information on Data Center High Availability, one can visit the below URL https://aws.amazon.com/answers/networking/aws-multiple-data-center-ha-network-connectivity/

● The traffic will flow primarily through the Internet

● Traffic will not flow due to the conflict in routes

Q20) You've setup a VPN connection between your on-premise data center and AWS. You need to know how the VPN connection will cost. Which of the below is a factor to be considered when looking at the costing for VPN connections? Choose 2 correct options.

● DataTransfer In

● DataTransfer Out

● Virtual Private Gateway transfer out

✓ VPN connection hours

Explanation:-By default all other options become invalid This is given in the AWS Documentationextension of your corporate datacenter. VPN Connection Pricing $0.05 per VPN Connection-hour $0.048 per VPN Connection-hour for connections to the Tokyo Region and Osaka-Local Region $0.065 per VPN Connection-hour for AWS GovCloud (US) Region If you choose to create a VPN Connection to your VPC using a Virtual Private Gateway, you are charged for each "VPN Connection-hour" that your VPN connection is provisioned and available. Each partial VPN Connection-hour consumed is billed as a full hour. You also incur standard AWS data transfer charges for all data transferred via the VPN Connection. If you no longer wish to be charged for a VPN Connection, you simply terminate your VPN Connection using the AWS Management Console, command line interface, or API. For more information on the pricing, one can visit the below URL https://aws.amazon.com/vpc/pricing/ https://aws.amazon.com/vpn/pricing/

Q21) Your company currently uses NAT instances to route traffic for Instances in private subnets. They need to convert these to NAT gateways to increase the amount of bandwidth required. They want to automate the provision. How can you accomplish this?

● Use AWS Inspector to replace the NAT instances with NAT gateways

● Use AWS Config to change the configuration of the NAT instance to a NAT gateway

✓ Use Cloudformation templates to replace the NAT instances with NAT gateways

Explanation:-This example is also given in the AWS Documentation Modifying your CloudFormation template to discontinue the use of NAT instances and consume NAT gateways is straightforward. You would: Allocate an Elastic IP address. However, it would not be directly assigned to an instance. Create a NAT gateway resource. Create a route to the Internet, but via the NAT gateway instead of going through a NAT instance. As in the code for NAT instances, this route would then be associated with the route table for the private subnets in the same Availability Zone. The updated example would look something like this: { ... "Resources" : { ... "NATGateway1EIP" : { "Type" : "AWS::EC2::EIP", "Properties" : { "Domain" : "vpc" } }, "NATGateway1" : { "Type" : "AWS::EC2::NatGateway", "DependsOn" : "VPCGatewayAttachment", "Properties" : { "AllocationId" : { "Fn::GetAtt" : ["NATGateway1EIP", "AllocationId"] }, "SubnetId" : { "Ref" : "PublicSubnetAZ1" } } }, "PrivateRoute1" : { "Type" : "AWS::EC2::Route", "Properties" : { "RouteTableId" : { "Ref" : "PrivateRouteTable1" }, "DestinationCidrBlock" : "0.0.0.0/0", "NatGatewayId" : { "Ref" : "NATGateway1" } } }, ... } ... } For more information on the using cloudformation templates for NAT gateways , one can visit the below URL https://aws.amazon.com/blogs/apn/taking-nat-to-the-next-level-in-aws-cloudformation-templates/

● Use Opswork to replace the NAT instances with NAT gateways

Q22) When configuring a Public VIF for AWS Direct Connect , which of the following is not required in the configuration

● VLAN ID

● Router Peer IP

● BGP ASN

✓ Virtual Private Gateway

Explanation:-For more information on creating a public VIF , one can visit the below URL https://docs.aws.amazon.com/directconnect/latest/UserGuide/create-vif.html

Q23) Your company currently has VPC's located in us-west and us-east. The company has an AWS Direct Connect connection in the US East region. They want to have the ability to extend the connection to us-west. They also need to minimize time and effort to have this in place. How can this be achieved?

✓ Make use of the Direct Connect gateway

Explanation:-The AWS Documentation mentions the following You can use an AWS Direct Connect gateway to connect your AWS Direct Connect connection over a private virtual interface to one or more VPCs in your account that are located in the same or different regions. You associate a Direct Connect gateway with the virtual private gateway for the VPC, and then create a private virtual interface for your AWS Direct Connect connection to the Direct Connect gateway. You can attach multiple private virtual interfaces to your Direct Connect gateway. For more information on the Direct Connect gateway , one can visit the below URL https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways.html

● Create another AWS Direct Connect connection in us-west
● Create a private VIF using the current connection
● Make use of an IPSec VPN

Q24) Your company has setup an application load balancer and various targets behind the ALB. But there are continuous problems at times wherein clients cannot connect to the ALB , because of the whitelisting that is required to be done by the IT Security department. What changes can be made to the architecture to alleviate this problem.

● Assign a public IP to the Application Load Balancer
● Assign an Elastic IP to the Application Load Balancer
✓ Place a Network Load balancer in front of the ALB

Explanation:-Since the IP of the Application Load balancer keeps on changing , the workaround is to have a Network Load balancer in front of the ALB. An elastic IP is then assigned to the Network Load balancer and in this way the IP address wont change. For more information on this scenario , one can visit the below URL https://aws.amazon.com/blogs/networking-and-content-delivery/using-static-ip-addresses-for-application-load-balancers/

● Place a Network Load balancer behind the ALB

Q25) You have a Cloudfront distribution that has an S3 bucket as the origin. There is a requirement to add Security headers to the response before it can be relayed back to the user. How can you achieve this?

● Change the Behaviour of the origin. Add a configuration for adding the security header.
✓ Create a Lambda function that will run on the edge

Explanation:-One of the AWS Blogs mentions the following [email protected] provides the ability to execute a Lambda function at an Amazon CloudFront Edge Location. This capability enables intelligent processing of HTTP requests at locations that are close (for the purposes of latency) to your customers. To get started, you simply upload your code (Lambda function written in Node.js) and pick one of the CloudFront behaviors associated with your distribution. All other options are incorrect since none of these will help meet the requirement For more information on adding security headers using [email protected] , one can visit the below URL https://aws.amazon.com/blogs/networking-and-content-delivery/adding-http-security-headers-using-lambdaedge-and-amazon-cloudfront/

● Make sure that the Viewer protocol is set to HTTPS
● Create an OAI for the Cloudfront distribution

Q26) Your company has setup a VPN connection between their on-premise infrastructure and AWS. They have multiple VPC's defined. They also need to ensure that all traffic flows through a transit VPC from their on-premise infrastructure. How would you architect the solution? Choose 2 answers from the options given below

● Create a VPN connection between the On-premise environment to all other VPC's
✓ Create a VPN connection between the transit VPC to all other VPC's

Explanation:-For more information on the transit VPC , one can visit the below URL https://docs.aws.amazon.com/aws-technical-content/latest/aws-vpc-connectivity-options/transit-vpc.html

✓ Create a VPN connection between the On-premise environment and the transit VPC

Explanation:-For more information on the transit VPC , one can visit the below URL https://docs.aws.amazon.com/aws-technical-content/latest/aws-vpc-connectivity-options/transit-vpc.html

● Create a VPC peering connection between the transit VPC and all other VPC's

Q27) Your company is planning on testing out Amazon workspaces for their account. They are going to allocate a set of workstations with static IP addresses for this purpose. They need to ensure that only these IP addresses have access to Amazon Workspaces. How can you achieve this?

● Specify the IP addresses in the NACL
● Specify the IP addresses in the Security Group
● Place a WAF in front of Amazon Workspaces
✓ Create an IP access control group

Explanation:-The AWS Documentation mentions the following An IP access control group acts as a virtual firewall that controls the IP addresses from which users are allowed to access their WorkSpaces. You can associate each IP access control group with one or more directories. You can associate up to 25 IP access control groups with each directory. For more information on restricting access , one can visit the below URL https://docs.aws.amazon.com/workspaces/latest/adminguide/amazon-workspaces-ip-access-control-groups.html

Q28) You have a RESTful service that is developed by your company. You want to provide secure access to this service to multiple clients in AWS. The service is hosted in a private subnet in one of your VPC's. How can you accomplish this?

● Create a VPC Endpoint gateway for your service
✓ Create a VPC Interface Endpoint that connects to the company's endpoint service

Explanation:-The AWS Documentation mentions the following An interface VPC endpoint (AWS PrivateLink) enables you to connect to services powered by AWS PrivateLink. These services include some AWS services, services hosted by other AWS accounts (referred to as endpoint services), and supported AWS Marketplace partner services. The interface endpoints are created directly inside of your VPC, using elastic network interfaces and IP addresses in your VPC's subnets. The service is now in your VPC, enabling connectivity to AWS services or AWS PrivateLink-powered service via private IP addresses. That means that VPC Security Groups can be used to manage access to the endpoints. Also, interface endpoint can be accessed from your premises via AWS Direct Connect. For more information on AWS private link , one can visit the below URL https://docs.aws.amazon.com/aws-technical-content/latest/aws-vpc-connectivity-options/aws-privatelink.html

● Create an application load balancer and provide the DNS name to your clients
● Create a network load balancer and provide the DNS name to your clients

Q29) You have a requirement of providing remote access to clients from their mobile devices and tablets. This is to access a service from inside a VPC. Which of the following would be part of the design?

● An AWS Managed Direct Connect connection
● A custom VPN server hosted on an EC2 Instance
✓ An AWS Client VPN
Explanation:-AWS now supports Client-to-site VPN Refer link : https://aws.amazon.com/about-aws/whats-new/2018/12/introducing-aws-client-vpn-to-securely-access-aws-and-on-premises-resources/. For more information on Remote Access connectivity , one can visit the below URL https://docs.aws.amazon.com/aws-technical-content/latest/aws-vpc-connectivity-options/software-remote-access-vpn-internal-user.html
● An AWS Managed Direct Connect gateway

Q30) Your company has setup an AWS Direct Connect connection along with a public VIF. There is a concern raised by the IT security department regarding the loopholes with a public VIF. Which of the following is a valid concern that could be raised by the security department?

✓ An EC2 Instance with a public IP has a chance of reaching you via the public VIF
Explanation:-For more information on the Reinvent video , one can visit the below URL https://www.youtube.com/watch?v=eNxPhHTN8gY
● An EC2 Instance with a private IP has a chance of reaching you via the public VIF
● Your VPC gets exposed via the public VIF
● Your VPC gets exposed to the Internet

Q31) Which of the following can be used to control how far your routes gets advertised when using AWS Direct Connect and a public VIF.

✓ Use BGP communities
Explanation:-This is also mentioned in the AWS Documentation BGP Communities AWS Direct Connect supports a range of BGP community tags to help control the scope (regional or global) and route preference of traffic. Scope BGP Communities You can apply BGP community tags on the public prefixes you advertise to Amazon to indicate how far to propagate your prefixes in the Amazon network—for the local AWS Region only, all regions within a continent, or all public regions. You can use the following BGP communities for your prefixes: 7224:9100—Local AWS Region 7224:9200—All AWS regions for a continent (for example, North America–wide) 7224:9300—Global (all public AWS Regions) All other options are invalid since you need to use BGP communities For more information on routing and BGP communities , one can visit the below URL https://docs.aws.amazon.com/directconnect/latest/UserGuide/routing-and-bgp.html
● Use BGP headers
● Use AS_PATH prepending
● Use MED

Q32) Your company is using a hosted private virtual interface from its parent AWS Account to the AWS Direct Connect Endpoint. You need to mention to IT management on what charges your company will acquire. Which of the following would you mention?

● The port hour charges
● The data transfer in
✓ The data transfer out via the private virtual interface to AWS Direct Connect Endpoint
Explanation:-The AWS Documentation currently mentions the following Data Transfer via AWS Direct Connect will be billed in the same month in which the usage occurred. If you have a hosted private virtual interface, you will only be charged for the data transferred out of that private virtual interface at the applicable Data Transfer rates. The account that owns the port will be charged the port-hour charges All other options are invalid since it is clearly mentioned what you get charged for in the AWS Documentation For more information on Direct Connect , one can visit the below URL https://aws.amazon.com/directconnect/faqs/
● The amount of hours used by the interface

Q33) You're working as a consultant for a company that has a three-tier application. The application layer of this architecture sends over 20Gbps of data per seconds during peak hours to and from Amazon S3. Currently, you're running two NAT gateways in two subnets to transfer the data from your private application layer to Amazon S3. You will also need to ensure that the instances receive software patches from a third-party repository, without leaving the AWS network. What architecture changes should be made, if any?

● Add another NAT gateway
✓ Add a VPC endpoint.
Explanation:-The AWS Documentation mentions the following A VPC endpoint enables you to privately connect your VPC to supported AWS services and VPC endpoint services powered by PrivateLink without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC do not require public IP addresses to communicate with resources in the service. Traffic between your VPC and the other service does not leave the Amazon network. For more information on VPC endpoints , one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-endpoints.html
● Add an Internet gateway for better throughput
● Add a VPN connection for better throughput

Q34) Your on-premise network has an IP address range of 10.55.0.0/16. You have been allocated an address range of 10.55.253.0/24 for the AWS Cloud. You need to design the VPC and ensure communication between the VPC and your on-premise network. You need to ensure proper set-up is configured at the customer end. How would you accomplish this? Choose 2 answers from the options given below

✓ Establish a VPN connection using your customer gateway. Ensure a route is present in your on-premise router to route traffic via the customer gateway.
Explanation:-Since the Address range assigned for the cloud is 10.55.253.0/24. This should be the address range assigned to the VPC Then use the customer gateway on your side to route traffic through the VPN tunnel. For more information on setting up a VPN connection, one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/SetUpVPNConnections.html
● Establish a VPN connection using your virtual private gateway. Ensure a route is present in your on-premise router to route traffic via the virtual private gateway.
● Setup a VPC with an address range of 10.55.0.0/16
✓ Setup a VPC with an address range of 10.55.253.0/24
Explanation:-Since the Address range assigned for the cloud is 10.55.253.0/24. This should be the address range assigned to the VPC Then use

the customer gateway on your side to route traffic through the VPN tunnel. For more information on setting up a VPN connection, one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/SetUpVPNConnections.html

Q35) You've setup a VPC with a couple of Instances that have public IP addresses. These EC2 Instances need to reach an external web server on port 443. The instances are unable to reach the web server. You have verified the following • An internet gateway is assigned to the VPC(10.0.0.0/16) • The route table has a route for 0.0.0.0/0 to the Internet gateway • The Security Groups allows Outbound Traffic on port 443 • The NACL allows Outbound Traffic on port 443 and Inbound Traffic for ephemeral ports Based on the above information what could be the underlying issue.

● You should not use the Internet gateway , instead use a NAT gateway for the routing of traffic
● The route table should have a route for 10.0.0.0/16 to the Internet gateway
● The route table is not having a route to the NAT gateway
✓ The external web server is blocking the requests

Explanation:-All of the settings are right for ensuring traffic can reach the external web server. In the end the issue could be at the web server end and it is blocking traffic. For more information on Amazon VPC, one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html

Q36) Your company is planning on setting up an AWS Direct Connect connection. Which of the following is not required for setting up the connection?

✓ Support for the router for IPSec

Explanation:-Below are the requirements for AWS Direct Connect • Your network must use single mode fiber with a 1000BASE-LX (1310nm) transceiver for 1 gigabit Ethernet, or a 10GBASE-LR (1310nm) transceiver for 10 gigabit Ethernet. • Auto-negotiation for the port must be disabled. Port speed and full-duplex mode must be configured manually. • 802.1Q VLAN encapsulation must be supported across the entire connection, including intermediate devices. • Your device must support Border Gateway Protocol (BGP) and BGP MD5 authentication. • (Optional) You can configure Bidirectional Forwarding Detection (BFD) on your network. Asynchronous BFD is automatically enabled for AWS Direct Connect virtual interfaces, but will not take effect until you configure it on your router. All other options are invalid because these are all key requirements for AWS Direct Connect For more information on AWS Direct Connect, one can visit the below URL https://docs.aws.amazon.com/directconnect/latest/UserGuide/dc-ug.pdf

● Support for the router for BGP
● Single mode fiber
● VLAN encapsulation

Q37) Your company is planning on using a Sub 1Gbps hosted connection. What is the final step that needs to be carried out before you can start creating a virtual interface using this connection?

● Create the hosted connection in the console
● Request for an AWS Direct Connect connection via AWS Support
✓ Accept the hosted connection in the console

Explanation:-This is mentioned in the AWS Documentation If you requested a sub-1G connection from your selected partner, they create a hosted connection for you (you cannot create it yourself). You must accept it in the AWS Direct Connect console before you can create a virtual interface. For more information on AWS Direct Connect, one can visit the below URL https://docs.aws.amazon.com/directconnect/latest/UserGuide/dc-ug.pdf

● Raise a support ticket to accept the hosted connection

Q38) You have an EC2 Instance that is located in a subnet mapped to an Availability zone. Due to a recent network redesign by your network architects the Instance needs to be moved to another availability zone. How can you achieve this?

✓ Create an AMI out of the EC2 Instance. Copy the AMI and launch a new Instance in the other Availability Zone

Explanation:-This is also given in the AWS Documentation How do I move my EC2 instance to another Availability Zone? Issue I want to move or copy my EC2 instance to another Availability Zone. How do I do that? Short Description To move an EC2 instance, create a new Amazon Machine Image (AMI) in the desired target Availability Zone, launch a new instance based on this image, and then reassign the Elastic IP address from the instance you are moving to the new image. All other options automatically become invalid because of this restriction For more information on moving EC2 Instances, one can visit the below URL https://aws.amazon.com/premiumsupport/knowledge-center/move-ec2-instance/

● Assign a new public IP address which pertains to the new subnet and then assign it to the Instance
● Create an ENI in the new subnet. Attach it to the Instance
● Assign a new private IP address which pertains to the new subnet and then assign it to the Instance

Q39) Your company has a set of EC2 Instances defined in a VPC. They need to monitor the traffic flowing into the Instances. They also need to monitor all the AWS API calls occurring on the EC2 Instances. Which of the following services can help fulfil this requirement?

● Amazon CloudWatch Logs and VPC Flow Logs
✓ AWS CloudTrail and VPC Flow Logs

Explanation:-The AWS Documentation mentions the following VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data is stored using Amazon CloudWatch Logs. After you've created a flow log, you can view and retrieve its data in Amazon CloudWatch Logs. AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs. For more information on VPC Flow logs, one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html

● AWS CloudTrail and CloudWatch Logs
● AWS CloudTrail and AWS Config

Q40) Your company currently has data in an on-premise location and they want to create a AWS Direct Connect connection to move this data to their AWS VPC, securely. You also need to access other AWS services and ensure confidentiality and integrity of the data in transit to your VPC. Choose 2 steps from the below options

✓ Setup a public VIF using the AWS Direct Connect connection

Explanation:-This is also given in the AWS Documentation Short Description A VPN that connects your office to your Amazon VPC over an AWS Direct Connect connection is likely to be faster and more secure than a VPN that connects to your VPC over the internet. Resolution Create an AWS Direct Connect connection. Configure a public virtual interface for the Direct Connect connection. In the Prefixes that you want to advertise, in the field for the virtual interface, enter the IPv4 CIDR destination addresses (separated by commas) where traffic should be routed to you over the virtual interface. In this case, add the public IP, as well as any network prefixes that you want to advertise. Your public virtual interface receives all the public IP addresses from AWS regions (except the AWS China region), including the public IP addresses of the VPN. To get the current list of prefixes advertised by AWS, download the JSON file containing AWS IP address ranges. For more information, see AWS IP Address Ranges. For more information on VPN over Direct Connect, one can visit the below URL https://aws.amazon.com/premiumsupport/knowledge-center/create-vpn-direct-connect/

● Setup a private VIF using the AWS Direct Connect connection

✓ Attach a virtual private gateway to the VPC

Explanation:-This is also given in the AWS Documentation Short Description A VPN that connects your office to your Amazon VPC over an AWS Direct Connect connection is likely to be faster and more secure than a VPN that connects to your VPC over the internet. Resolution Create an AWS Direct Connect connection. Configure a public virtual interface for the Direct Connect connection. In the Prefixes that you want to advertise, in the field for the virtual interface, enter the IPv4 CIDR destination addresses (separated by commas) where traffic should be routed to you over the virtual interface. In this case, add the public IP, as well as any network prefixes that you want to advertise. Your public virtual interface receives all the public IP addresses from AWS regions (except the AWS China region), including the public IP addresses of the VPN. To get the current list of prefixes advertised by AWS, download the JSON file containing AWS IP address ranges. For more information, see AWS IP Address Ranges. For more information on VPN over Direct Connect, one can visit the below URL https://aws.amazon.com/premiumsupport/knowledge-center/create-vpn-direct-connect/

● Create a IPSec tunnel between the customer gateway and the virtual private gateway

---

Q41) You have configured a classic load balancer in the public subnet with EC2 instances behind them. You are sending an HTTP request using the DNS name as the destination, but you are not getting the response from the underlying instances. Which of the following are checks you should carry out? Choose 2 answers from the options given below

✓ Ensure the load balancer is created in the public subnet

Explanation:-These checks are also given in the AWS Documentation Troubleshoot a Classic Load Balancer: Client Connectivity If your Internet-facing load balancer in a VPC is not responding to requests, check for the following: Your Internet-facing load balancer is attached to a private subnet Verify that you specified public subnets for your load balancer. A public subnet has a route to the Internet Gateway for your virtual private cloud (VPC). A security group or network ACL does not allow traffic The security group for the load balancer and any network ACLs for the load balancer subnets must allow inbound traffic from the clients and outbound traffic to the clients on the listener ports.

● Ensure the load balancer is created in the private subnet

● Ensure the Security group for the load balancer accepts traffic on port 80 from10.0.0.0/16

✓ Ensure the Security group for the load balancer accepts traffic on port 80 from0.0.0.0/0

Explanation:-These checks are also given in the AWS Documentation Troubleshoot a Classic Load Balancer: Client Connectivity If your Internet-facing load balancer in a VPC is not responding to requests, check for the following: Your Internet-facing load balancer is attached to a private subnet Verify that you specified public subnets for your load balancer. A public subnet has a route to the Internet Gateway for your virtual private cloud (VPC). A security group or network ACL does not allow traffic The security group for the load balancer and any network ACLs for the load balancer subnets must allow inbound traffic from the clients and outbound traffic to the clients on the listener ports.

Q42) You're trying to do some housekeeping and delete some unwanted interface. You try to delete an interface manually that has the following information { "Status": "in-use", ... "Description": "VPC Endpoint Interface vpce-08233123488812123", "NetworkInterfaceId": "eni-c8fbc27e", "VpcId": "vpc-1a2b3c4d", "PrivateIpAddresses": [{ "PrivateDnsName": "ip-20-0-2-227.ec2.internal", "Primary": true, "PrivateIpAddress": "20.0.2.227" }], "RequesterManaged": true, ... } But you are not able to delete the interface. What is the reason as to why you cannot delete the interface?

✓ It's because it is a requester managed interface

Explanation:-The AWS Documentation mentions the following A requester-managed network interface is a network interface that an AWS service creates in your VPC. This network interface can represent an instance for another service, such as an Amazon RDS instance, or it can enable you to access another service or resource, such as an AWS PrivateLink service, or an Amazon ECS task. You cannot modify or detach a requester-managed network interface. If you delete the resource that the network interface represents, the AWS service detaches and deletes the network interface for you. For more information on requester managed interfaces, one can visit the below URL https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/requester-managed-eni.html

● It's because it has a private DNS name attached

● It's because it has a private IP address attached

● It's because its attached to a VPC

Q43) An EC2 Instance has been setup in AWS. A software was successfully downloaded and installed on the EC2 Instance. This software uses IPv6 for communication. After the software was installed , and you were trying to access the software via IPv6 on port 80, you were not able to do so. What needs to be done to alleviate this issue?

✓ Add an inbound rule to your security group that allows inbound traffic on port 80 for ::/0.

Explanation:-Since the application works on IPv6, you need to ensure that the port is open for all Ipv6 addresses as ::/0. For more information on security groups, one can visit the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html

● Add an internet gateway for the instance.

● Add an inbound rule to your security group that allows inbound traffic on port 80 for 0.0.0.0/0.

● Add an egress-only internet gateway.

Q44) You have a collection of assets stored in an S3 bucket. You want to enable users across the world to access these assets with the least latency. The users must also access the distribution via your company domain name. How can you achieve this? Choose 2 answers from the options given below.

✓ Create a web based distribution in Cloudfront

Explanation:-This is also given in the AWS Documentation Routing Traffic to an Amazon CloudFront Web Distribution by Using Your Domain Name If you want to speed up delivery of your web content, you can use Amazon CloudFront, the AWS content delivery network (CDN). CloudFront can deliver your entire website—including dynamic, static, streaming, and interactive content—by using a global network of edge locations. Requests for your content are automatically routed to the edge location that gives your users the lowest latency. Note You can route traffic to a CloudFront distribution only for public hosted zones. To use CloudFront to distribute your content, you create a web distribution and specify settings such as the

Amazon S3 bucket or HTTP server that you want CloudFront to get your content from, whether you want only selected users to have access to your content, and whether you want to require users to use HTTPS. When you create a web distribution, CloudFront assigns a domain name to the distribution, such as d111111abcdef8.cloudfront.net. You can use this domain name in the URLs for your content, for example: http://d111111abcdef8.cloudfront.net/logo.jpg Alternatively, you might prefer to use your own domain name in URLs, for example: http://example.com/logo.jpg If you want to use your own domain name, use Amazon Route 53 to create an alias record that points to your CloudFront distribution. An alias record is a Route 53 extension to DNS. It's similar to a CNAME record, but you can create an alias record both for the root domain, such as example.com, and for subdomains, such as www.example.com. (You can create CNAME records only for subdomains.) When Route 53 receives a DNS query that matches the name and type of an alias record, Route 53 responds with the domain name that is associated with your distribution. For more information on routing to Cloudfront, one can visit the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfront-distribution.html

● Create an application load balancer and point it to your S3 bucket

✓ Create a resource record in a hosted zone and create an ALIAS record

Explanation:-This is also given in the AWS Documentation Routing Traffic to an Amazon CloudFront Web Distribution by Using Your Domain Name If you want to speed up delivery of your web content, you can use Amazon CloudFront, the AWS content delivery network (CDN). CloudFront can deliver your entire website—including dynamic, static, streaming, and interactive content—by using a global network of edge locations. Requests for your content are automatically routed to the edge location that gives your users the lowest latency. Note You can route traffic to a CloudFront distribution only for public hosted zones. To use CloudFront to distribute your content, you create a web distribution and specify settings such as the Amazon S3 bucket or HTTP server that you want CloudFront to get your content from, whether you want only selected users to have access to your content, and whether you want to require users to use HTTPS. When you create a web distribution, CloudFront assigns a domain name to the distribution, such as d111111abcdef8.cloudfront.net. You can use this domain name in the URLs for your content, for example: http://d111111abcdef8.cloudfront.net/logo.jpg Alternatively, you might prefer to use your own domain name in URLs, for example: http://example.com/logo.jpg If you want to use your own domain name, use Amazon Route 53 to create an alias record that points to your CloudFront distribution. An alias record is a Route 53 extension to DNS. It's similar to a CNAME record, but you can create an alias record both for the root domain, such as example.com, and for subdomains, such as www.example.com. (You can create CNAME records only for subdomains.) When Route 53 receives a DNS query that matches the name and type of an alias record, Route 53 responds with the domain name that is associated with your distribution. For more information on routing to Cloudfront, one can visit the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfront-distribution.html

● Create a resource record in a hosted zone and create a PTR record

Q45) A windows machine in one VPC needs to join the AD domain in another VPC. VPC Peering has been established. But the domain join is not working. What is the other step that needs to be followed to ensure that the AD domain join can work as intended

● Change the VPC peering connection to a VPNconnection

● Change the VPC peering connection to a Direct Connect connection

✓ Ensure the security groups for AD hosted instance has the right rules for relevant instances.

Explanation:-In addition to VPC peering and setting the right route tables , the security groups for the AD EC2 instance needs to ensure the right rules are put in place for allowing incoming traffic. For more information on allowing ingress traffic for AD, please visit the following url https://docs.aws.amazon.com/quickstart/latest/active-directory-ds/ingress.html

● Ensure that the AD is placed in a public subnet

Q46) Which of the following is a key prerequisite for creating an AWS Managed Microsoft AD directory? Choose 2 answers from the options given below

✓ A VPC with 2 subnets

Explanation:-The AWS Documentation mentions the following To create an AWS Managed Microsoft AD directory, you need a VPC with the following: • At least two subnets. Each of the subnets must be in a different Availability Zone. • The following ports must be open between the two subnets that you deploy your directory into. This is necessary to allow the domain controllers that AWS Directory Service creates for you to communicate with each other. A security group will be created and attached to your directory to enable communication between the domain controllers. o TCP/UDP 53 - DNS o TCP/UDP 88 - Kerberos authentication o UDP 123 - NTP o TCP 135 - RPC o UDP 137-138 - Netlogon o TCP 139 - Netlogon o TCP/UDP 389 - LDAP o TCP/UDP 445 - SMB o TCP 636 - LDAPS (LDAP over TLS/SSL) o TCP 873 - Rsync o TCP 3268 - Global Catalog o TCP/UDP 1024-65535 - Ephemeral ports for RPC • The VPC must have default hardware tenancy. • You cannot create a AWS Managed Microsoft AD in a VPC using addresses in the 198.19.0.0/16 address space. • AWS Directory Service does not support using Network Address Translation (NAT) with Active Directory. Using NAT can result in replication errors. For more information on the pre-requisites, please visit the following url https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_getting_started_prereqs.html

● Usage of a NAT Instance in the VPC

✓ Opening of several ports including port 53

Explanation:-The AWS Documentation mentions the following To create an AWS Managed Microsoft AD directory, you need a VPC with the following: • At least two subnets. Each of the subnets must be in a different Availability Zone. • The following ports must be open between the two subnets that you deploy your directory into. This is necessary to allow the domain controllers that AWS Directory Service creates for you to communicate with each other. A security group will be created and attached to your directory to enable communication between the domain controllers. o TCP/UDP 53 - DNS o TCP/UDP 88 - Kerberos authentication o UDP 123 - NTP o TCP 135 - RPC o UDP 137-138 - Netlogon o TCP 139 - Netlogon o TCP/UDP 389 - LDAP o TCP/UDP 445 - SMB o TCP 636 - LDAPS (LDAP over TLS/SSL) o TCP 873 - Rsync o TCP 3268 - Global Catalog o TCP/UDP 1024-65535 - Ephemeral ports for RPC • The VPC must have default hardware tenancy. • You cannot create a AWS Managed Microsoft AD in a VPC using addresses in the 198.19.0.0/16 address space. • AWS Directory Service does not support using Network Address Translation (NAT) with Active Directory. Using NAT can result in replication errors. For more information on the pre-requisites, please visit the following url https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_getting_started_prereqs.html

● A NAT gateway in the public subnet

Q47) Your company currently has a set of EC2 Instances hosted in a VPC. The IT Security department wants to find out the details of the traffic that had caused an issue on one of the instances. What can you do to zero in on the IP addresses which are receiving a flurry of requests.

✓ Use VPC Flow logs to get the IP addresses accessing the EC2 Instances

Explanation:-With VPC Flow logs you can get the list of IP addresses which are hitting (or had hit) the Instances in your VPC. You can then use the information in the logs to see which external IP addresses are sending a flurry of requests which could be the potential threat for causing the issue. For more information on VPC Flow Logs, please visit the following url https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html

● Use AWS Cloud trail to get the IP addresses accessing the EC2 Instances

● Use AWS Config to get the IP addresses accessing the EC2 Instances
● Use AWS Trusted Advisor to get the IP addresses accessing the EC2 Instances

Q48) Your company has a set of EC2 Instances that are placed behind an ELB. Some of the applications hosted on these instances communicate via a legacy protocol. There is a security mandate that all traffic between the client and the EC2 Instances need to be secure. How would you accomplish this?

● Use an Application Load balancer and terminate the SSL connection at the ELB
● Use a Classic Load balancer and terminate the SSL connection at the ELB
● Use an Application Load balancer and terminate the SSL connection at the EC2 Instances
✓ Use a Classic Load balancer and terminate the SSL connection at the EC2 Instances
Explanation:-Since there are applications which work on legacy protocols, you need to ensure that the ELB can be used at the network layer as well and hence you should choose the Classic ELB. Since the traffic needs to be secure till the EC2 Instances , the SSL termination should occur on the Ec2 Instances. For more information on HTTPS listeners for classic load balancers, please refer to below URL
https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-https-load-balancers.html

Q49) Your company is planning on setting up a Direct Connect connection to AWS. But they don’t require or have the facility to accommodate a 1Gbps connection. How can they achieve a sub 1G connection? Choose 2 answers from the options given below.

✓ If they have a parent AWS Account which can accommodate a 1G connection , look at having a Hosted Virtual Interface
Explanation:-Below are the options as given in the AWS Documentation Hosted virtual interfaces (VIF) Hosted virtual interfaces (VIF) can connect to public resources or a VPC in the same way as standard VIFs, except that the account that owns the hosted VIF is different from the connection owner. Bandwidth is shared across all virtual interfaces on the parent connection. Hosted connections allow an APN partner to create a Direct Connect sub-1G connection for you, allocating dedicated bandwidth for that connection rather than having multiple VIFs on the same parent connection competing for bandwidth. For more information on the Direct Connect types , please visit the following URL
https://aws.amazon.com/premiumsupport/knowledge-center/direct-connect-types/
● If they have a parent AWS Account which can accommodate a 1G connection , look at having a Hosted Connection
● They can consider contacting an AWS Partner for a Hosted Virtual Interface
✓ They can consider contacting an AWS Partner for a Hosted Connection
Explanation:-Below are the options as given in the AWS Documentation Hosted virtual interfaces (VIF) Hosted virtual interfaces (VIF) can connect to public resources or a VPC in the same way as standard VIFs, except that the account that owns the hosted VIF is different from the connection owner. Bandwidth is shared across all virtual interfaces on the parent connection. Hosted connections allow an APN partner to create a Direct Connect sub-1G connection for you, allocating dedicated bandwidth for that connection rather than having multiple VIFs on the same parent connection competing for bandwidth. For more information on the Direct Connect types , please visit the following URL
https://aws.amazon.com/premiumsupport/knowledge-center/direct-connect-types/

Q50) A Company currently uses the NetFlow software to monitor and get the details of the traffic that flows between systems in their On-premise network. They want to have the same ability when they start moving their servers to AWS. Which of the following service can help them meet this requirement.

● AWS Cloudwatch logs
✓ AWS VPC Flow Logs
Explanation:-NetFlow is a network protocol developed by Cisco for collecting IP traffic information and monitoring network traffic. The AWS Documentation mentions the following VPC Flow Logs are similar to scheduled NetFlow/sFlow/IPFIX reports. Flow logs collect the source and destination IP, source and destination ports, protocol, packet counts, and ALLOW or DENY action for a particular VPC, subnet, or ENI. They are currently collected and sent as a report every 10 minutes. For more information on VPC Flow logs , please visit the following URL
https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html
● AWS Cloudwatch metrics
● AWS Config

Q51) A company has an application that needs to be moved to an AWS VPC network. This application is based on multicast and needs to be moved with the least amount of effort. What can be done to fulfil this requirement?

● Create EC2 Instances in the subnet and then migrate the application on to the EC2 Instance.
● The application needs to be changed to support unicast before moving it to AWS.
✓ Consider creating an overlay network between EC2 Instances and then port the application.
Explanation:-Currently Amazon VPC service doesn’t presently permit multicast or broadcast traffic. In the event that you have an application that uses multicast to function, you can leverage GRE tunnels to create a mesh VPN overlay network between your Amazon EC2 instances Option A is incorrect because you need to setup an overlay network first.
https://aws.amazon.com/marketplace/pp/B071RMCZ1X
● Consider enabling encryption on the underlying EBS volumes which will be used to support the EC2 Instance

Q52) A company has acquired another company. Both companies have their presence in AWS and in the same region, US-East. There is a requirement to ensure EC2 inside VPC A of the parent company and EC2 of VPC B of the parent company can communicate with each other. Also, ensure EC2 inside VPC B of the parent company and EC2 inside VPC C of the acquired company can communicate with each other. CIDR of VPC of each VPC are as follows VPC A: 10.9.0.0/16 VPC B: 10.11.0.0/16 VPC C: 172.16.0.0/16 How can you accomplish this Architecture?

✓ Create a VPC Peering connection between VPC A and VPC B. Create another VPC peering connection between VPC B and VPC C
Explanation:-Requirement: traffic from EC2 of VPC A and EC2 of VPC B of the same company can communicate Traffic from EC2 of VPC B of current company and EC2 of VPC C of the acquired company can communicate CIDR of each Each VPC is different Both companies have resources in the US-East region. It talks about to create 2 VPC. create VPC peering between VPC A and VPC B Create VPC peering between VPC B and VPC C This is possible as CIDR of all VPC is different and are in the same region. For more information on transit networks, please refer to the below URL https://aws.amazon.com/answers/networking/aws-global-transit-network/
● Create a VPC Peering connection between VPC A and VPC C. Create another VPC peering connection between VPC B and VPC C
● Create a VPC Peering connection between VPC A and VPC B. Create a VPN connection between VPC B and VPC C
● Create a VPC Peering connection between VPC A and VPC C. Create a VPN connection between VPC A and VPC B

Q53) You currently have setup a VPN configuration from your on-premise location to AWS. Your AWS VPC has a CIDR of 10.0.0.0/16 and a subnet of 10.0.1.0/24. Your On-premise location has a network CIDR block of 10.0.37.0/24 and 10.1.38.0/24. The traffic is being dropped when it is being sent from the subnet instances to your on-premise location. Why could be the most probable reason in this case?

● You have not set Enhanced Networking on the Instances

✓ There is an overlap in prefixes

Explanation:-Such an example is given in the AWS Documentation Connections with Your Local Network and Other VPCs You can optionally set up a connection between your VPC and your corporate or home network. If you have an IPv4 address prefix in your VPC that overlaps with one of your networks' prefixes, any traffic to the network's prefix is dropped. For example, let's say that you have the following: A VPC with CIDR block 10.0.0.0/16 A subnet in that VPC with CIDR block 10.0.1.0/24 Instances running in that subnet with IP addresses 10.0.1.4 and 10.0.1.5 On-premises host networks using CIDR blocks 10.0.37.0/24 and 10.1.38.0/24 When those instances in the VPC try to talk to hosts in the 10.0.37.0/24 address space, the traffic is dropped because 10.0.37.0/24 is part of the larger prefix assigned to the VPC (10.0.0.0/16). The instances can talk to hosts in the 10.1.38.0/24 space because that block isn't part of 10.0.0.0/16. For more information on VPC and Subnets , please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

● The “Do not fragment” is set in the IP header

● The MTU is not set to 9001

Q54) Your company is planning on creating a Direct Connect connection and also have a VPN as a backup connection. Which of the following must be done to ensure that the AWS Direct connect connection is the preferred path?

✓ Ensure that prefixes are advertised the same on both connections

Explanation:-By default , AWS will choose AWS Direct Connect. In order to ensure architecture for proper failover the AWS Documentation mentions the following points To configure the hardware VPN as a backup for your Direct Connect connection: • Be sure that you use the same virtual private gateway for both Direct Connect and the VPN connection to the VPC. • If you are configuring a Border Gateway Protocol (BGP) VPN, advertise the same prefix for Direct Connect and the VPN. • If you are configuring a static VPN, add the same static prefixes to the VPN connection that you are announcing with the Direct Connect virtual interface. • If you are advertising the same routes toward the AWS VPC, the Direct Connect path is always be preferred, regardless of AS path prepending. For more information on these points , please refer to the below URL https://aws.amazon.com/premiumsupport/knowledge-center/configure-vpn-backup-dx/

● Ensure that the longest prefix is advertised on AWS Direct connect

● Ensure that AS_PATH prepending is configured on AWS Direct Connect

● Ensure that the shortest prefix is advertised on AWS Direct connect

Q55) Your company is planning on trying out AWS Workspaces for 100 users. They want to have a standalone managed directory service along with AWS workspaces. Which of the following would be the ideal option which will have a least administrative overhead and also be cost effective.

● Deploy an AD domain server in a VPC and configure AWS Workspace to use the newly created AD Domain server

● Choose an AD connector to use along with AWS Workspaces

✓ Choose Simple AD to use along with AWS Workspaces

Explanation:-The AWS Documentation mentions the following Amazon WorkSpaces uses directories to store and manage information for your WorkSpaces and users. For your directory, you can choose from Simple AD, AD Connector, or AWS Directory Service for Microsoft Active Directory (Enterprise Edition), also known as Microsoft AD. In addition, you can establish a trust relationship between your Microsoft AD directory and your on-premises domain. For more information on Simple AD with AWS workspaces , please refer to the below URL https://docs.aws.amazon.com/workspaces/latest/adminguide/launch-workspace-simple-ad.html

● Choose AWS Directory Service to use along with AWS Workspaces

Q56) You are setting up a VPN software on an EC2 Instance which will be used for VPN connections. Which of the following is an important aspect that should be set on the EC2 Instance?

✓ Disable source destination check on the Amazon EC2 instance.

Explanation:-An example is also given in the AWS Documentation To launch an EC2 VPN instance 1.Launch an Amazon Linux instance in a VPC public subnet and do the following: a)Assign the VPN instance a static private IP address. This is not required, but it makes setting up the config files easier. In this example, use 10.0.0.5. b)Allocate a VPC EIP and associate an EIP to your VPN instance. In this example, use EIP1 to represent the public EIP address used to connect into your VPC. 2. Disable Source/Dest checking on your EC2 instance. a)Right-click the instance and selecting Change Source/Dest. Check. b)Click Yes, Disable.

● Enable source destination check on the Amazon EC2 instance.

● Enable route propagation in a Virtual Private Cloud (VPC) subnet route table.

● Enable enhanced networking mode on the Amazon EC2 instance.

Q57) Your company is planning on setting up an AWS Direct connect connection to an AWS VPC. They want to achieve maximum fault tolerance , have maximum bandwidth at all times. How can this be achieved?

✓ One Virtual Private gateway Two AWS Direct Connect Locations Two Customer gateways

Explanation:-For more information on high network connectivity , please refer to the below URL https://aws.amazon.com/answers/networking/aws-multiple-data-center-ha-network-connectivity/

● Two Virtual Private gateway Two AWS Direct Connect Locations One Customer gateway

● Two Virtual Private gateway One AWS Direct Connect Location One Customer gateway

● One Virtual Private gateway One AWS Direct Connect Location One VPN connection Two Customer gateways

Q58) Your company currently has a VPC hosted in AWS. There is a private hosted zone in place for the instances in this VPC. You need your On-premise servers to be able to resolve the AD DNS requests for instances in the VPC. You need to do this with the least amount of effort. What steps would you. Choose 2 answers from the options given below.

✓ Setup a Simple AD in AWS.

Explanation:-The AWS Documentation mentions the following Simple AD forwards DNS requests to the IP address of the Amazon-provided DNS servers for your VPC. These DNS servers will resolve names configured in your Route 53 private hosted zones. By pointing your on-premises computers to your Simple AD, you can now resolve DNS requests to the private hosted zone. For more information on Simple AD , please refer to the below URL https://docs.aws.amazon.com/directoryservice/latest/admin-guide/simple_ad_dns.html

● Setup an Active Directory Domain Controller in the AWS VPC

✓ Make your On-premise servers point to the Simple AD

Explanation:-The AWS Documentation mentions the following Simple AD forwards DNS requests to the IP address of the Amazon-provided DNS servers for your VPC. These DNS servers will resolve names configured in your Route 53 private hosted zones. By pointing your on-premises computers to your Simple AD, you can now resolve DNS requests to the private hosted zone. For more information on Simple AD , please refer to the below URL https://docs.aws.amazon.com/directoryservice/latest/admin-guide/simple_ad_dns.html

● Make your On-premise servers point to the new Domain Controller

Q59) You've currently configured health checks in Route 53. These health checks are being used for 2 of your on-premise web servers. The health checks are not working as desired. The health checks are continually failing. Which of the following could be a possible reason?

● Ensure that the Security groups on the Instances are allowing Inbound Traffic

● Ensure that the NACL's on the Subnets are allowing Inbound Traffic

✓ Ensure that the Firewall on your On-premise environment is allowing Inbound Traffic

Explanation:-The AWS Documentation mentions the following When Route 53 checks the health of an endpoint, it sends an HTTP, HTTPS, or TCP request to the IP address and port that you specified when you created the health check. For a health check to succeed, your router and firewall rules must allow inbound traffic from the IP addresses that the Route 53 health checkers use. For more information on Route 53 health checks , please refer to the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-router-firewall-rules.html

● This is not possible. You cannot enable health checks for non-AWS resources

Q60) You are currently configuring Route 53 routing policies. You want to create a record set for a group of Web servers in your AWS VPC. When a user requests for the resource record , they should be able to access any of the web servers defined in the VPC. Which of the following resource record would you create?

● Simple

● Weighted

● Failover

✓ Multivalue answer

Explanation:-The AWS Documentation mentions the following Multivalue answer routing lets you configure Amazon Route 53 to return multiple values, such as IP addresses for your web servers, in response to DNS queries. You can specify multiple values for almost any record, but multivalue answer routing also lets you check the health of each resource, so Route 53 returns only values for healthy resources. It's not a substitute for a load balancer, but the ability to return multiple health-checkable IP addresses is a way to use DNS to improve availability and load balancing. For more information on Routing policy's , please refer to the below URL https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

Q61) Your team is planning on creating a set of instances in a VPC. They need to ensure high network performance for the underlying instances and enhanced communication between the instances. Which of the following steps would you take. Choose 2 answers from the options given below

✓ Enable Enhanced Networking for the underlying Instances

Explanation:-The AWS Documentation mentions the following A cluster placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information on Placement groups , please refer to the below URL https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

● Set the MTU for the Instances to 1500

● Create the Instances in separate Availability Zones and put them in a cluster placement Group

✓ Create the Instances in the same Availability Zones and put them in a cluster placement Group

Explanation:-The AWS Documentation mentions the following A cluster placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information on Placement groups , please refer to the below URL https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

Q62) You are trying to send packets from an EC2 Instance to an on-premise server. The transmission is happening over the Internet. You have set Jumbo frames due to the size of the packets being sent. But the packets are being dropped. What needs to be done to ensure that the packets don't get dropped?

● Ensure that the MTU is set to 9001

● Ensure that the "Do Not Fragment" flag is set in the IP header

✓ Ensure that the "Do Not Fragment" flag is not set in the IP header

Explanation:-The AWS Documentation mentions the following Jumbo frames allow more than 1500 bytes of data by increasing the payload size per packet, and thus increasing the percentage of the packet that is not packet overhead. Fewer packets are needed to send the same amount of usable data. However, outside of a given AWS region (EC2-Classic), a single VPC, or a VPC peering connection, you will experience a maximum path of 1500 MTU. VPN connections and traffic sent over an Internet gateway are limited to 1500 MTU. If packets are over 1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header. For more information on Network MTU , please refer to the below URL https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/network_mtu.html

● Enable Enhanced Networking on the Instance

Q63) You're hosting an NGINX web server running on port 80 on an EC2 Instance. Users are not able to access the server running on port 80. Which of the following could be an issue

● The Security Group does not allow outbound traffic on port 80

✓ The NACL don't allow outbound traffic on ephemeral ports

Explanation:-When a connection is established on a client , you need to ensure that outbound traffic is enabled on any ephemeral ports for the client. This is also given in the AWS Documentation Ephemeral Ports The example network ACL in the preceding section uses an ephemeral port range of 32768-65535. However, you might want to use a different range for your network ACLs depending on the type of client that you're using or with which you're communicating. The client that initiates the request chooses the ephemeral port range. The range varies depending on the client's operating system. Many Linux kernels (including the Amazon Linux kernel) use ports 32768-61000. Requests originating from Elastic Load Balancing use ports 1024-65535. Windows operating systems through Windows Server 2003 use ports 1025-5000. Windows Server 2008 and later versions use ports 49152-65535. A NAT gateway uses ports 1024-65535. For example, if a request comes into a web server in your VPC from a Windows XP client on the Internet, your network ACL must have an outbound rule to enable traffic destined for ports 1025-5000. For more information on NACL 's, please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

● The NACL don't allow inbound traffic on ephemeral ports
● The Security Group does not allow inbound traffic on ephemeral ports

Q64) Your company is planning on moving its files from its on-premise location onto S3. The services hosted in the on-premise environment would need low latency access to these files. How can you achieve this?

● Create a VPN connection which would allow the services on-premise to access S3

✓ Create a Direct Connect connection along with a Public VIF

Explanation:-This is also given in the AWS Documentation To connect to AWS public endpoints, such as an Amazon Elastic Compute Cloud (Amazon EC2) or Amazon Simple Storage Service (Amazon S3), with dedicated network performance, use a public virtual interface. For more information on Public and Private Virtual Interfaces , please refer to the below URL https://aws.amazon.com/premiumsupport/knowledge-center/public-private-interface-dx/

● Create a Direct Connect connection along with a Private VIF

● Create a VPN connection along with a VPC endpoint

Q65) You have a set of Instances in your VPC that communicate over the IPv6 protocol. You need to ensure that traffic can flow from the Instances to the Internet but not vice versa. How ca you achieve this.

● Change the Internet gateway to only allow outbound traffic for IPv6

● Change the Security Groups to not allow Inbound Traffic on the Instances

● Change the NACL's to not allow Inbound Traffic on the Instances

✓ Use an Egress only Internet gateway

Explanation:-This is also given in the AWS Documentation IPv6 addresses are globally unique, and are therefore public by default. If you want your instance to be able to access the Internet, but you want to prevent resources on the Internet from initiating communication with your instance, you can use an egress-only Internet gateway. To do this, create an egress-only Internet gateway in your VPC, and then add a route to your route table that points all IPv6 traffic (::/0) or a specific range of IPv6 address to the egress-only Internet gateway. IPv6 traffic in the subnet that's associated with the route table is routed to the egress-only Internet gateway. For more information on Egress only Internet gateway , please refer to the below URL https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/egress-only-internet-gateway.html