AWS_ANS-C00 Exam

Volume: 315 Questions

Question No:1
A customer you configured a Direct Connect (DX) connection for would like to send traffic from a VPC in the US-East-1 region over their corporate network backbone for auditing and on to a VPC in US-West-1 region. Via BGP, the US-East-1 region virtual private gateway (VGW) is advertising an AS number of AS 7224, the corporate Cisco router is advertising an AS number of AS 64511, the VGW in US-West-1 is advertising an AS number of AS 7224. The VGW at US-West-1 is rejecting the traffic originating from US-East-1. Which option below would best resolve the issue?

A. Create a software VPN connection between the two VPCs

B. Use VPC Peering to connect the two VPCs

C. Create a Bastion host on the corporate network and configure security groups to allow inbound connections from US-East-1 and outbound connections to US-West-1

D. Use AS-Override on the corporate router

Answer: D

Explanation: In a BGP environment, if a router sees the same AS number it has in an AS_PATH, it will reject the traffic as the same AS number suggests a loop. In this scenario, enabling AS-Override on the corporate router would advertise AS 64511 AS 64511 to the VGW at US-West-1 instead of AS 64511 AS 7224, allowing the traffic to be accepted. VPC Peering or a software VPN connection would not meet the requirement for routing the traffic through the corporate router. A Bastion host would not be a practical solution for this use case.
Reference: https://www.youtube.com/watch?v=Qep11X1r1QA#t=30m30s

Question No:2
A client has a Classic Load Balancer named “client-lb” configured to distribute incoming requests across 10 instances in the us-east-1a AZ and 5 instances in the us-west-1a AZ. The client has a requirement for the incoming requests to be distributed equally across all 15 instances. Using the AWS CLI, which command below could be used to meet this requirement?

A. aws elb create-load-balancer --load-balancer-name client-lb --listeners "Protocol=HTTP,LoadBalancerPort=80,InstanceProtocol=HTTP,InstancePort=80" --availability-zones us-east-1a us-west-1a

B. aws elb modify-load-balancer-attributes --load-balancer-name client-lb

--load-balancer-attributes "{\"ConnectionDraining\":{\"Enabled\":true,\"Timeout\":300}}"

C. aws elb modify-load-balancer-attributes --load-balancer-name client-lb --load-balancer-attributes "{\"CrossZoneLoadBalancing\":{\"Enabled\":true}}"

D. aws elb create-load-balancer --load-balancer-name client-lb --listeners "Protocol=HTTP,LoadBalancerPort=443,InstanceProtocol=HTTP,InstancePort=443"
--availability-zones us-east-1a us-west-1a

Answer: C

Explanation: Cross-Zone Load Balancing distributes incoming requests evenly across all enabled Availability Zones. This reduces the need to have equal numbers of instances in all Availability Zones and improves the ability of an application to deal with the loss of instances. The "modify-load-balancer" can be used in conjunction with the --load-balancer-attributes option and "{\"CrossZoneLoadBalancing\":{\"Enabled\":true}}" variables to enable Cross-Zone Load Balancing on an existing Classic Load Balancer.
Reference:
http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/enable-disable-crosszone-lb.html#enable-cross-zone

Question No:3
Your company was recently suffered a security breach. Since then you have implemented more restrictive policies both on-premises and in AWS. In addition to the NACL, IAM, and security group policies you have put in place, you would like to be alerted if there are 5 rejected SSH attempts to your EC2 instance with an ENI within any given hour. You will configure a CloudWatch Metric Filter and Alarm for a Flow Log to achieve this requirement. Which of the following should NOT be part of your Filter Pattern?

A. destport="22"

B. action="REJECT"

C. action="ALARM"

D. protocol="6"

Answer: C

Explanation: The alarm for a metric filter in CloudWatch is configured separately from the filter pattern. The “action” value in a filter pattern refers to what happened to the network traffic. In this example, you are looking for traffic that has been rejected, so you should filter for “REJECT”.

For SSH connections, the "destport" value should be "22" (the default SSH port) and the "protocol" value should be "6" (TCP).
Reference:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html#flow-logs-cw-alarm-example

Question No:4
You are the network engineer responsible for a VPN connection between a VPC's virtual private gateway (VGW) and your on-premises VPN device (customer gateway). You are attempting to establish the IPsec tunnel between the two devices. Your customer gateway is on the AWS list of "Customer Gateway Devices We've Tested". NAT-T is not being used. You have configured both the VGW and customer gateway and are now attempting to establish the IPsec tunnel between them. Phase 1 (IKE exchange) is successful. Phase 2 (IPsec phase) fails. You have tried multiple times with the same results. You have ruled out incompatible IKE versions or encryption methods as the cause. Which option below should you check first?

A. Verify UDP port 500 is not blocked

B. Verify you can ping the customer gateway

C. Verify protocol 50 is not blocked

D. Verify UDP port 4500 is not blocked

Answer: C

Explanation: Protocol 50 being blocked is the only option listed that would impact phase 2 but not phase 1. If UDP port 500 was blocked, phase 1 would have failed. UDP port 4500 would only be required if NAT-T was used. Ping would be used to verify network connectivity to the customer gateway which is confirmed by the fact phase 1 was successful.
Reference:
https://aws.amazon.com/premiumsupport/knowledge-center/vpn-tunnel-phase-2-ipsec/

Question No:5
Your account has 3 VPCs deployed in the same region.
The IPv4 address ranges of the VPCs are:
VPC A 172.22.0.0/16
VPC B 10.3.0.0/16
VPC C 10.4.0.0/16
VPC B and VPC C are already connected through VPC peering connection (PCX) named pcx-abcxyz. VPC A has resources that VPC B and VPC C must access. Which of the options

below best achieves the objective of allowing VPC B and VPC C to access VPC A?

A. Connect VPC A to pcx-abcxyz

B. Add static routes to VPC A with Targets of 10.3.0.0/16 and 10.4.0.0/16 and a Destination of pcx-abcxyz

C. Create separate VPC Peering connections between VPC A & VPC B and VPC A & VPC C

D. Connect VPC A to VPC B using VPC Peering and allow VPC C to access VPC A through VPC B.

Answer: C

Explanation: Creating separate VPC Peering connections between each VPC is required for Peering 3 VPCs together. Transitive routing is not supported in VPC Peering so VPC C could not reach VPC A through VPC B. A new VPC Peering connection is required for connecting additional VPCs, so connecting VPC A to pcx-abcxyz is not a viable option. Adding a static route to VPC A that points to pcx-abcxyz would not work as pcx-abcxyz is not connected to VPC A.
Reference:
http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html#three-vpcs-full-access

Question No:6
Your customer will be storing critical data in S3 buckets. In addition to a highly restrictive bucket policy, they would like to encrypt the data prior to loading it to the buckets and avoid storing master keys off premises. Which of the following options would it be best to leverage to encrypt the data given the requirements?

A. SSE-S3

B. SSE-KMS

C. KMS–Managed CMK

D. Client-Side Master Key

Answer: D

Explanation: Given the requirement to encrypt the data prior to loading it to the S3 buckets, server side solutions like SSE-S3 and SSE-KMS would not meet the requirements. While a KMS-Managed CMK would allow for client-side encryption, it would not keep master key out of AWS. Using a Client-Side Master Key would allow for client-side encryption and keep the master

key out of AWS.
Reference:
http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingClientSideEncryption.html

Question No:7
Your company has a customer gateway (CGW) that you will be using to connect to two separate VPCs via VPN. There are overlapping private IP address ranges on the two VPCs. You need to mitigate the potential issues that can occur when overlapping IP address ranges exist on a routing table. Which solution best meets the objective?

A. Peer the VPCs and connect the CGW to one VPC only

B. Change the AS number of the CGW

C. Use VRF on the CGW

D. Use DX to connect the VPCs and connect the CGW to one VPC only

Answer: C

Explanation: 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.
Reference: https://aws.amazon.com/articles/5458758371599914

Question No:8
Your customer is currently manually deploying EC2 instances to a VPC and provisioning them to serve various roles in their web application. The customer has a limited operations staff and developers are spending a significant amount of time installing packages & frameworks and configuring software. The customer would like to automate the installation of packages & frameworks and software configuration as well as scale EC2 deployments based on traffic to allow the developers to focus on development as opposed to infrastructure. There is a preference to minimize cost for any solution implemented. Which solution should you recommend?

A. Kinesis

B. CloudWatch

C. OpsWorks for Chef Automate

D. OpsWorks Stacks

Answer: D

Explanation: OpsWorks stacks allows for automation of package installation, installation of programming language and frameworks, and software configuration. It also allows you to scale based on changing traffic levels and is offered with no additional charge for use with EC2 (you only pay for the resources created). OpsWorks for Chef Automate has costs associated with the number of nodes connected to the Chef server and the underlying instance running the Chef Server. CloudWatch is a monitoring solution. Kinesis is used to collect and load streaming data.
Reference: https://aws.amazon.com/opsworks/

Question No:9
A customer would like to optimize the performance of their web application by routing inbound traffic to api.customersite.net to Compute Optimized EC2 instances and inbound traffic to mobile.customersite.net to Memory Optimized EC2 instances. Which solution below would be best to implement for this customer?

A. Enable X-Forwarded For on the web servers and use a Classic Load Balancer

B. Use an Application Load Balancer with path-based routing rules to forward the traffic to the correct instances

C. Configure proxy servers to forward the traffic to the correct instances

D. Use an Application Load Balancer with host-based routing rules to forward the traffic to the correct instances

Answer: D

Explanation: Application Load Balancers can be used with host-based routing rules to route traffic by subdomain names (i.e. X.customersite.net and Y.customersite.net) . Path-based routing takes into account the path in the URI (i.e. /2017/july/blogs) and would not be helpful in routing traffic based on subdomains. Proxy servers could be used to achieve the desired functionality, but are less efficient. X-Forwarded For would not be helpful in forwarding traffic to specific instances.
Reference:
https://aws.amazon.com/blogs/aws/new-host-based-routing-support-for-aws-application-load-balancers/

Question No:10

Your company was delivering training content to employees and select customers using an S3 bucket configured as a static website. The content consisted of static HTML pages and image files. Your company would like to enhance the content by adding video that will be streamed using the RTMP protocol and increase security by restricting access to the content to only employees. Previously, anyone with the S3 URLs could access the content. You have already decided on CloudFront as the CDN and denied public access to the S3 URLs. Which of the following options will best meet the requirement of restricting access to employees?

A. Create an origin access identity, give the origin access identity read permissions to the S3 bucket, and use signed URLs

B. Use signed cookies and signed URLs

C. Create an origin access identity, give the origin access identity read permissions to the S3 bucket, and use signed cookies

D. Create a bucket policy giving read access to requests originating on your corporate network and use signed cookies

Answer: A

Explanation: An origin access identity is a special user associated with CloudFront that can be used to grant access to content within a distribution. In this scenario, that user would need read access to the S3 bucket. Signed cookies are not supported with RTMP content, so signed URLs should be used instead.
Reference:
http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/PrivateContent.htm
l

Question No:11
Your customer has a number of EC2 instances deployed in an auto scaling group. The auto scaling policies scale up and down based on CPU utilization. The demand on the instances is highly variable, so auto scaling has proven to be quite useful in the application. After the latest spike in traffic, an EC2 instance that served as the master node for a Hadoop cluster was terminated when CPU utilization went back down. The customer needs to prevent this from occurring again. Which solution should you recommend?

A. Increase the cooldown period for the auto scaling group

B. Disable auto scaling and “right size” the EC2 instances to handle the highest traffic levels expected

C. Increase the Desired Capacity for the auto scaling group

D. Enable scale in protection on the EC2 instance in question

Answer: D

Explanation: Enabling scale in protection on the instance would stop it from being terminated when auto scaling policies scale down. Disabling auto scaling would be overkill and is not desirable given that it serves a logical purpose in this scenario. Increasing the cooldown period or Desired Capacity may prevent the instance from being terminated depending on other variables, but it is not a surefire way to prevent termination long term.
Reference: https://aws.amazon.com/blogs/aws/new-instance-protection-for-auto-scaling/

Question No:12
You are the network engineer responsible a VPN connection between a VPC’s virtual private gateway (VGW) and your on-premises VPN device (customer gateway). You are attempting to establish the IPsec tunnel between the two devices. Your customer gateway is on the AWS list of “Customer Gateway Devices We've Tested”. NAT-T is not being used. You have configured both the VGW and customer gateway and are now attempting to establish the IPsec tunnel between them. Phase 1 of the IKE exchange fails. You have ruled out incompatible IKE versions or encryption methods as the cause. Which option below could be creating the failure?

A. Your ISP is blocking outbound connections to UDP port 4500

B. Your ISP is blocking outbound connections to TCP port 4500

C. Your ISP is blocking inbound connections to TCP port 500

D. Your ISP is blocking inbound connections to UDP port 500

Answer: D

Explanation: UDP port 500 needs to be open between a VGW and customer gateway for an IPsec tunnel to be established. Port 4500 is only required if using NAT-T.
Reference:
https://aws.amazon.com/premiumsupport/knowledge-center/vpn-tunnel-phase-1-ike/

Question No:13
You are using OpsWorks Stacks to manage instances that act as application servers deployed in a single VPC. The average usage is consistent at all hours of the day with occasional bursts in traffic that can be handled by 1 additional instance. The maximum amount of instances that

have ever been used in this application is 6. You are looking to minimize cost without sacrificing performance. What combination of instances should you deploy?

A. 24/7 instances only

B. time-based and load-based instances

C. 24/7, time-based, and load-based instances

D. 24/7 and load-based instances

Answer: D

Explanation: While it is generally recommended to deploy a mixture of 24/7, time-based, and load-based instances, in this scenario, there is no need for time-based instances as traffic is consistent with bursts that can be handled on an as-needed basis by load-based instances.
Reference:
http://docs.aws.amazon.com/opsworks/latest/userguide/best-practices-autoscale.html

Question No:14
You have Linux EC2 instances deployed in 2 VPCs. VPC A has an IP address range of 10.20.30.0/24. VPC B has an IP address range of 10.20.31.0/24. The DNS servers supplied by the DHCP options for both subnets reside in VPC A. VPC B cannot ping the hostnames of instances in VPC A. Instances in VPC A can ping the hostnames of other instances in VPC A. VPC B can ping the IPv4 addresses of the DNS servers. Other than the DNS servers, all the instances have dynamically assigned IPv4 addresses. Which of the following is most likely the cause?

A. IGMP is disabled on the instances in VPC A

B. UDP port 67 outbound is blocked on VPC B

C. ICMP is disabled on the instances running in VPC A

D. UDP port 53 outbound is blocked on VPC B

Answer: D

Explanation: Based on the symptoms, it seems that instances in VPC B cannot resolve the hostnames of the instances in VPC A. This would most likely be caused by access to a DNS server being blocked. DNS queries occur on UDP (or in certain circumstances TCP) port 53. UDP Port 67 is used by DHCP servers. Since ping is working under certain circumstances, ICMP is

not disabled. IGMP is not used for ping.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-dns.html

Question No:15
Your on-premises datacenter is connected to your VPC using a DX connection. You have multiple Windows EC2 instances deployed within a private subnet in your VPC. You have configured the instances security group rules to allow for RDP access from your datacenter so you can access the instances for administration. You would like to further harden the configuration and mitigate risk of MITM attacks. Which of the following should you implement to mitigate risk to MITM attacks?

A. SSL encryption of the RDP connection

B. EV SSL Certificates

C. X.509 certificates

D. Change the default RDP port

Answer: C

Explanation: By default, Windows RDP uses untrusted, self-signed certificates. X.509 certificates can be used to ensure you are communicating directly with the Windows instance as opposed to a “man-in-the-middle”. RDP connections establish an underlying SSL/TLS connection, so manual SSL encryption of the RDP connection would be unnecessary. EV SSL Certificates are used for websites. Changing the default RDP port would be “security by obscurity” at best, but would not inherently mitigate MITM attacks.
Reference: https://d0.awsstatic.com/whitepapers/Security/AWS_Security_Best_Practices.pdf

Question No:16
When an Amazon Virtual Private Cloud (VPC) is created, Amazon automatically creates a set of DHCP options. One of the options is AmazonProvidedDNS which maps an Amazon provided DNS server to an IPv4 address that is reserved at the base of the VPC’s IPv4 network range, plus two. Given this information, what would be the IP address for the Amazon provided DNS server in a VPC with a network range of 10.20.30.0/24?

A. 10.20.30.2

B. 10.20.30.1

C. 10.20.32.1