Free Questions for ANS-C01 by certscare

Shared by Valenzuela on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Which solution will meet these requirements MOST cost-effectively?

Options:

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

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

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

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

Answer:

D

Explanation:

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

Question 2

Question Type: MultipleChoice

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

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

Which solution will meet these requirements?

Options:

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

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

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

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

Answer:

A

Question 3

Question Type: MultipleChoice

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

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

What action will resolve the availability problem?

Options:

A- Create a new subnet using a VPC secondary IPv6 CIDR, and associate an IPv6 CIDR. Include the new subnet in the Auto Scaling group.

B- Create a new subnet using a VPC secondary IPv4 CIDR, and associate an IPv6 CIDR. Include the new subnet in the Auto Scaling group.

C- Resize the IPv6 CIDR on each of the existing subnets. Modify the Auto Scaling group maximum number of instances.

D- Add a secondary IPv4 CIDR to the Amazon VPC. Assign secondary IPv4 address space to each of the existing subnets.

Answer:

B

Question 4

Question Type: MultipleChoice

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

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

Protocol: TCP

Port: 80 inbound, nothing outbound

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

Protocol: TCP

Port: 80 inbound, nothing outbound

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

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

Options:

A- Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 80

B- Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 1024-65535

C- Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 80

D- Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 1024-65535

Answer:

D

Explanation:

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

Question 5

Question Type: MultipleChoice

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

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

How should you use AWS services to meet these requirements?

Options:

A- Register the IP addresses of the service hosts as "A" records with latency-based routing policy in Amazon Route 53, and set a Route 53 health check for these hosts.

B- Set the Elastic Load Balancing (ELB) load balancer in front of the hosts of the service, and register the ELB name of the main service host as an ALIAS record with a latency-based routing policy in Route 53.

C- Set Amazon CloudFront in front of the host of the service, and register the CloudFront name of the main service as an ALIAS record in Route 53.

D- Set the Amazon API gateway in front of the service, and register the API gateway name of the main service as an ALIAS record in Route 53.

Answer:

B

Question 6

Question Type: MultipleChoice

A customer has set up multiple VPCs for Dev, Test, Prod, and Management. You need to set up AWS Direct Connect to enable data flow from on-premises to each VPC. The customer has monitoring software running in the Management VPC that collects metrics from the instances in all the other VPCs. Due to budget requirements, data transfer charges should be kept at minimum.

Which design should be recommended?

Options:

A- Create a total of four private VIFs, one for each VPC owned by the customer, and route traffic between VPCs using the Direct Connect link.

B- Create a private VIF to the Management VPC, and peer this VPC to all other VPCs.

C- Create a private VIF to the Management VPC, and peer this VPC to all other VPCs, enable source/destination NAT in the Management VPC.

D- Create a total of four private VIFs, and enable VPC peering between all VPCs.

Answer:

D

Explanation:

- creating VPC peering is free of charge - traffic costs ~0.01/GB for VPC peering (IN + OUT) and ~0.02/GB for direct connect (OUT only). As the communication involved in monitoring will never have IN == OUT, then 0.01 * (IN + OUT) will always be lower the 0.02 * OUT, ergo VPC peering will be cheaper

Question 7

Question Type: MultipleChoice

Your security team implements a host-based firewall on all of your Amazon Elastic Compute Cloud (EC2) instances to block all outgoing traffic. Exceptions must be requested for each specific requirement. Until you request a new rule, you cannot access the instance metadata service. Which firewall rule should you request to be added to your instances to allow instance metadata access?

Options:

A- Inbound; Protocol tcp; Source [Instance's EIP]; Destination 169.254.169.254

B- Inbound; Protocol tcp; Destination 169.254.169.254; Destination port 80

C- Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 80

D- Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 443

Answer:

C

Explanation:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html

To view all categories of instance metadata from within a running instance, use the following URI. http://169.254.169.254/latest/meta-data/

Question 8

Question Type: MultipleChoice

Your organization has a newly installed 1-Gbps AWS Direct Connect connection. You order the cross-connect from the Direct Connect location provider to the port on your router in the same facility. To enable the use of your first virtual interface, your router must be configured appropriately.

What are the minimum requirements for your router?

Options:

A- 1-Gbps Multi Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.

B- 1-Gbps Single Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.

C- IPsec Parameters, Pre-Shared key, Peer IP Address, BGP Session with MD5

D- BGP Session with MD5, 802.1Q VLAN, Route-Map, Prefix List, IPsec encrypted GRE Tunnel

Answer:

B