

ANS-C01 Dumps

AWS Certified Advanced Networking Specialty Exam

<https://www.certleader.com/ANS-C01-dumps.html>



NEW QUESTION 1

A company has an AWS Site-to-Site VPN connection between its existing VPC and on-premises network. The default DHCP options set is associated with the VPC. The company has an application that is running on an Amazon Linux 2 Amazon EC2 instance in the VPC. The application must retrieve an Amazon RDS database secret that is stored in AWS Secrets Manager through a private VPC endpoint. An on-premises application provides internal RESTful API service that can be reached by URL (<https://api.example.internal>). Two on-premises Windows DNS servers provide internal DNS resolution. The application on the EC2 instance needs to call the internal API service that is deployed in the on-premises environment. When the application on the EC2 instance attempts to call the internal API service by referring to the hostname that is assigned to the service, the call fails. When a network engineer tests the API service call from the same EC2 instance by using the API service's IP address, the call is successful. What should the network engineer do to resolve this issue and prevent the same problem from affecting other resources in the VPC?

- A. Create a new DHCP options set that specifies the on-premises Windows DNS server
- B. Associate the new DHCP options set with the existing VP
- C. Reboot the Amazon Linux 2 EC2 instance.
- D. Create an Amazon Route 53 Resolver rul
- E. Associate the rule with the VP
- F. Configure the rule to forward DNS queries to the on-premises Windows DNS servers if the domain name matches example.internal.
- G. Modify the local host file in the Amazon Linux 2 EC2 instance in the VPMap the service domain name (api.example.internal) to the IP address of the internal API service.
- H. Modify the local /etc/resolv.conf file in the Amazon Linux 2 EC2 instance in the VP
- I. Change the IP addresses of the name servers in the file to the IP addresses of the company's on-premises Windows DNS servers.

Answer: B

Explanation:

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

NEW QUESTION 2

A company has developed an application on AWS that will track inventory levels of vending machines and initiate the restocking process automatically. The company plans to integrate this application with vending machines and deploy the vending machines in several markets around the world. The application resides in a VPC in the us-east-1 Region. The application consists of an Amazon Elastic Container Service (Amazon ECS) cluster behind an Application Load Balancer (ALB). The communication from the vending machines to the application happens over HTTPS.

The company is planning to use an AWS Global Accelerator accelerator and configure static IP addresses of the accelerator in the vending machines for application endpoint access. The application must be accessible only through the accelerator and not through a direct connection over the internet to the ALB endpoint.

Which solution will meet these requirements?

- A. Configure the ALB in a private subnet of the VP
- B. Attach an internet gateway without adding routes in the subnet route tables to point to the internet gatewa
- C. Configure the accelerator with endpoint groups that include the ALB endpoint
- D. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- E. Configure the ALB in a private subnet of the VP
- F. Configure the accelerator with endpoint groups that include the ALB endpoint
- G. Configure the ALB's security group to only allow inbound traffic from the internet on the ALB listener port.
- H. Configure the ALB in a public subnet of the VPAttach an internet gatewa
- I. Add routes in the subnet route tables to point to the internet gatewa
- J. Configure the accelerator with endpoint groups that include the ALB endpoint
- K. Configure the ALB's security group to only allow inbound traffic from the accelerator's IP addresses on the ALB listener port.
- L. Configure the ALB in a private subnet of the VP
- M. Attach an internet gatewa
- N. Add routes in the subnet route tables to point to the internet gatewa
- O. Configure the accelerator with endpoint groups that include the ALB endpoint
- P. Configure the ALB's security group to only allow inbound trafficfrom the accelerator's IP addresses on the ALB listener port.

Answer: A

Explanation:

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

NEW QUESTION 3

A network engineer must develop an AWS CloudFormation template that can create a virtual private gateway, a customer gateway, a VPN connection, and static routes in a route table. During testing of the template, the network engineer notes that the CloudFormation template has encountered an error and is rolling back. What should the network engineer do to resolve the error?

- A. Change the order of resource creation in the CloudFormation template.
- B. Add the DependsOn attribute to the resource declaration for the virtual private gatewa
- C. Specify the route table entry resource.
- D. Add a wait condition in the template to wait for the creation of the virtual private gateway.
- E. Add the DependsOn attribute to the resource declaration for the route table entr
- F. Specify the virtual private gateway resource.

Answer: D

NEW QUESTION 4

A data analytics company has a 100-node high performance computing (HPC) cluster. The HPC cluster is for parallel data processing and is hosted in a VPC in the AWS Cloud. As part of the data processing workflow, the HPC cluster needs to perform several DNS queries to resolve and connect to Amazon RDS databases, Amazon S3 buckets, and on-premises data stores that are accessible through AWS Direct Connect. The HPC cluster can increase in size by five to seven times during the company's peak event at the end of the year.

The company is using two Amazon EC2 instances as primary DNS servers for the VPC. The EC2 instances are configured to forward queries to the default VPC resolver for Amazon Route 53 hosted domains and to the on-premises DNS servers for other on-premises hosted domain names. The company notices job failures and finds that DNS queries from the HPC cluster nodes failed when the nodes tried to resolve RDS and S3 bucket endpoints.

Which architectural change should a network engineer implement to provide the DNS service in the MOST scalable way?

- A. Scale out the DNS service by adding two additional EC2 instances in the VP
- B. Reconfigure half of the HPC cluster nodes to use these new DNS server
- C. Plan to scale out by adding additional EC2instance-based DNS servers in the future as the HPC cluster size grows.
- D. Scale up the existing EC2 instances that the company is using as DNS server
- E. Change the instance size to the largest possible instance size to accommodate the current DNS load and theanticipated load in the future.
- F. Create Route 53 Resolver outbound endpoint
- G. Create Route 53 Resolver rules to forward queries to on-premises DNS servers for on premises hosted domain name
- H. Reconfigure the HPC cluster nodes to use the default VPC resolver instead of the EC2 instance-based DNS server
- I. Terminate the EC2 instances.
- J. Create Route 53 Resolver inbound endpoint
- K. Create rules on the on-premises DNS servers to forward queries to the default VPC resolve
- L. Reconfigure the HPC cluster nodes to forward all DNS queries to the on-premises DNS server
- M. Terminate the EC2 instances.

Answer: C

NEW QUESTION 5

A company has an AWS Direct Connect connection between its on-premises data center in the United States (US) and workloads in the us-east-1 Region. The connection uses a transit VIF to connect the data center to a transit gateway in us-east-1.

The company is opening a new office in Europe with a new on-premises data center in England. A Direct Connect connection will connect the new data center with some workloads that are running in a single VPC in the eu-west-2 Region. The company needs to connect the US data center and us-east-1 with the Europe data center and eu-west-2. A network engineer must establish full connectivity between the data centers and Regions with the lowest possible latency.

How should the network engineer design the network architecture to meet these requirements?

- A. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VI
- B. Associate the transit gateway in us-east-1 with the same Direct Connect gatewa
- C. Enable SiteLink for the transit VIF and the private VIF.
- D. Connect the VPC in eu-west-2 to a new transit gatewa
- E. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VI
- F. Associate the transit gateway in us-east-1 with the same Direct Connect gatewa
- G. Enable SiteLink for both transit VIF
- H. Peer the two transit gateways.
- I. Connect the VPC in eu-west-2 to a new transit gatewa
- J. Connect the Europe data center to the new transit gateway by using a Direct Connect gateway and a new transit VI
- K. Create a new Direct Connect gatewa
- L. Associate the transit gateway in us-east-1 with the new Direct Connect gatewa
- M. Enable SiteLink for both transit VIF
- N. Peer the two transit gateways.
- O. Connect the VPC in eu-west-2 with the Europe data center by using a Direct Connect gateway and a private VI
- P. Create a new Direct Connect gatewa
- Q. Associate the transit gateway in us-east-1 with the new Direct Connect gatewa
- R. Enable SiteLink for the transit VIF and the private VIF.

Answer: C

NEW QUESTION 6

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

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

- A. Create a central egress VPC that has private NAT gateway
- B. Connect all the VPCs to the central egress VPC by using AWS Transit Gatewa
- C. Use the private NAT gateways to connect to Amazon S3 and Systems Manager by using private IP addresses.
- D. Create a central shared services VP
- E. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
- F. Ensure that private DNS is turned of
- G. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
- H. Create an Amazon Route 53 forwarding rule for each interface VPC endpoint
- I. Associate the forwarding rules with all the VPC
- J. Forward DNS queries to the interface VPC endpoints in the shared services VPC.
- K. Create a central shared services VPIIn the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
- L. Ensure that private DNS is turned of
- M. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
- N. Create an Amazon Route 53 private hosted zone with a full service endpoint name for Amazon S3 and Systems Manage
- O. Associate the private hosted zones with all the VPC
- P. Create an alias record in each private hosted zone with the full AWS service endpoint pointing to the interface VPC endpoint in the shared services VPC.
- Q. Create a central shared services VP
- R. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
- S. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
- T. Ensure that private DNS is turned on for the interface VPC endpoints and that the transit gateway is created with DNS support turned on.

Answer: B

Explanation:

Interface VPC endpoints enable private connectivity between VPCs and supported AWS services without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection². Interface VPC endpoints are powered by AWS PrivateLink, a technology that enables private access to AWS services². Amazon S3 and AWS Systems Manager support interface VPC endpoint²s. By turning off private DNS, the interface VPC endpoints can be accessed by using their private IP addresses². By using Amazon Route 53 forwarding rules, DNS queries can be resolved to the interface VPC endpoints in the shared services VPC³.

NEW QUESTION 7

A real estate company is building an internal application so that real estate agents can upload photos and videos of various properties. The application will store these photos and videos in an Amazon S3 bucket as objects and will use Amazon DynamoDB to store corresponding metadata. The S3 bucket will be configured to publish all PUT events for new object uploads to an Amazon Simple Queue Service (Amazon SQS) queue.

A compute cluster of Amazon EC2 instances will poll the SQS queue to find out about newly uploaded objects. The cluster will retrieve new objects, perform proprietary image and video recognition and classification update metadata in DynamoDB and replace the objects with new watermarked objects. The company does not want public IP addresses on the EC2 instances.

Which networking design solution will meet these requirements MOST cost-effectively as application usage increases?

- A. Place the EC2 instances in a public subnet
- B. Disable the Auto-assign Public IP option while launching the EC2 instance
- C. Create an internet gateway
- D. Attach the internet gateway to the VPC
- E. In the public subnet's route table, add a default route that points to the internet gateway.
- F. Place the EC2 instances in a private subnet
- G. Create a NAT gateway in a public subnet in the same Availability Zone
- H. Create an internet gateway
- I. Attach the internet gateway to the VPC
- J. In the public subnet's route table, add a default route that points to the internet gateway
- K. Place the EC2 instances in a private subnet
- L. Create an interface VPC endpoint for Amazon SQS
- M. Create gateway VPC endpoints for Amazon S3 and DynamoDB.
- N. Place the EC2 instances in a private subnet
- O. Create a gateway VPC endpoint for Amazon SQS. Create interface VPC endpoints for Amazon S3 and DynamoDB.

Answer: C

NEW QUESTION 8

A company's network engineer is designing a hybrid DNS solution for an AWS Cloud workload. Individual teams want to manage their own DNS hostnames for their applications in their development environment. The solution must integrate the application-specific hostnames with the centrally managed DNS hostnames from the on-premises network and must provide bidirectional name resolution. The solution also must minimize management overhead.

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

- A. Use an Amazon Route 53 Resolver inbound endpoint.
- B. Modify the DHCP options set by setting a custom DNS server value.
- C. Use an Amazon Route 53 Resolver outbound endpoint.
- D. Create DNS proxy servers.
- E. Create Amazon Route 53 private hosted zones.
- F. Set up a zone transfer between Amazon Route 53 and the on-premises DNS.

Answer: ABE

NEW QUESTION 9

A network engineer needs to set up an Amazon EC2 Auto Scaling group to run a Linux-based network appliance in a highly available architecture. The network engineer is configuring the new launch template for the Auto Scaling group.

In addition to the primary network interface the network appliance requires a second network interface that will be used exclusively by the application to exchange traffic with hosts over the internet. The company has set up a Bring Your Own IP (BYOIP) pool that includes an Elastic IP address that should be used as the public IP address for the second network interface.

How can the network engineer implement the required architecture?

- A. Configure the two network interfaces in the launch template
- B. Define the primary network interface to be created in one of the private subnets
- C. For the second network interface, select one of the public subnets
- D. Choose the BYOIP pool ID as the source of public IP addresses.
- E. Configure the primary network interface in a private subnet in the launch template
- F. Use the user data option to run a cloud-init script after boot to attach the second network interface from a subnet with auto-assign public IP addressing enabled.
- G. Create an AWS Lambda function to run as a lifecycle hook of the Auto Scaling group when an instance is launched
- H. In the Lambda function, assign a network interface to an AWS Global Accelerator endpoint.
- I. During creation of the Auto Scaling group, select subnets for the primary network interface
- J. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

Answer: D

Explanation:

During creation of the Auto Scaling group, select subnets for the primary network interface. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

This solution meets all of the requirements stated in the question. The primary network interface can be configured in a private subnet during creation of the Auto Scaling group. The user data option can be used to run a cloud-init script that will allocate a second network interface and associate an Elastic IP address from the BYOIP pool with it.

NEW QUESTION 10

A software-as-a-service (SaaS) provider hosts its solution on Amazon EC2 instances within a VPC in the AWS Cloud. All of the provider's customers also have their environments in the AWS Cloud.

A recent design meeting revealed that the customers have IP address overlap with the provider's AWS deployment. The customers have stated that they will not share their internal IP addresses and that they do not want to connect to the provider's SaaS service over the internet.

Which combination of steps is part of a solution that meets these requirements? (Choose two.)

- A. Deploy the SaaS service endpoint behind a Network Load Balancer.
- B. Configure an endpoint service, and grant the customers permission to create a connection to the endpoint service.
- C. Deploy the SaaS service endpoint behind an Application Load Balancer.
- D. Configure a VPC peering connection to the customer VPC
- E. Route traffic through NAT gateways.
- F. Deploy an AWS Transit Gateway, and connect the SaaS VPC to i
- G. Share the transit gateway with the customer
- H. Configure routing on the transit gateway.

Answer: AB

Explanation:

NLB for creating the private link which solves the overlapping IP address issue and the SaaS service endpoint behind it. (the SaaS endpoint could be an ALB)
<https://aws.amazon.com/about-aws/whats-new/2021/09/application-load-balancer-aws-privatelink-static-ip>

NEW QUESTION 10

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

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

- Protocol: TCP
- Port: 80 inbound, nothing outbound

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

- Protocol: TCP
- Port: 80 inbound, nothing outbound

When you try to browse to the web server, you receive no response. Which additional step should you take to receive a successful response?

- A. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 80
- B. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 1024-65535
- C. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 80
- D. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 1024-65535

Answer: D

Explanation:

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

NEW QUESTION 12

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

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

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

- A. Set up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional traffic load from remote employees and the additional applicatio
- B. Create a link aggregation group (LAG).
- C. Deploy an AWS Site-to-Site VPN connection to the application VP
- D. Configure the on-premises routing for the remote employees to connect to the Site-to-Site VPN connection.
- E. Deploy Amazon Workspaces into the application VPI
- F. Replace the existing 1 Gbps Direct Connect connection with two new 2 Gbps Direct Connect hosted connection
- G. Create an AWS Client VPN endpoint in the application VP
- H. Instruct the remote employees to connect to the Client VPN endpoint.

Answer: A

Explanation:

Setting up a new 1 Gbps Direct Connect dedicated connection to accommodate the additional trafficload from remote employees and the additional application would provide more bandwidth and lower latency than a VPN connection over the public internet1. Creating a link aggregation group (LAG) with the existing and new Direct Connect connections would provide resiliency and redundancy for the AWS connectivity2.

NEW QUESTION 16

A company is building its website on AWS in a single VPC. The VPC has public subnets and private subnets in two Availability Zones. The website has static content such as images. The company is using Amazon S3 to store the content.

The company has deployed a fleet of Amazon EC2 instances as web servers in a private subnet. The EC2 instances are in an Auto Scaling group behind an Application Load Balancer. The EC2 instances will serve traffic, and they must pull content from an S3 bucket to render the webpages. The company is using AWS Direct Connect with a public VIF for on-premises connectivity to the S3 bucket.

A network engineer notices that traffic between the EC2 instances and Amazon S3 is routing through a NAT gateway. As traffic increases, the company's costs are increasing. The network engineer needs to change the connectivity to reduce the NAT gateway costs that result from the traffic between the EC2 instances and Amazon S3.

Which solution will meet these requirements?

- A. Create a Direct Connect private VI
- B. Migrate the traffic from the public VIF to the private VIF.
- C. Create an AWS Site-to-Site VPN tunnel over the existing public VIF.
- D. Implement interface VPC endpoints for Amazon S3. Update the VPC route table.
- E. Implement gateway VPC endpoints for Amazon S3. Update the VPC route table.

Answer: D

NEW QUESTION 17

A global delivery company is modernizing its fleet management system. The company has several business units. Each business unit designs and maintains applications that are hosted in its own AWS account in separate application VPCs in the same AWS Region. Each business unit's applications are designed to get data from a central shared services VPC.

The company wants the network connectivity architecture to provide granular security controls. The architecture also must be able to scale as more business units consume data from the central shared services VPC in the future.

Which solution will meet these requirements in the MOST secure manner?

- A. Create a central transit gatewa
- B. Create a VPC attachment to each application VP
- C. Provide full mesh connectivity between all the VPCs by using the transit gateway.
- D. Create VPC peering connections between the central shared services VPC and each application VPC in each business unit's AWS account.
- E. Create VPC endpoint services powered by AWS PrivateLink in the central shared services VPCCreate VPC endpoints in each application VPC.
- F. Create a central transit VPC with a VPN appliance from AWS Marketplac
- G. Create a VPN attachment from each VPC to the transit VP
- H. Provide full mesh connectivity among all the VPCs.

Answer: C

Explanation:

Option C provides a secure and scalable solution using VPC endpoint services powered by AWS PrivateLink. AWS PrivateLink enables private connectivity between VPCs and services without exposing the data to the public internet or using a VPN connection. By creating VPC endpoints in each application VPC, the company can securely access the central shared services VPC without the need for complex network configurations. Furthermore, PrivateLink supports cross-account connectivity, which makes it a scalable solution as more business units consume data from the central shared services VPC in the future.

NEW QUESTION 18

A company is using Amazon Route 53 Resolver DNS Firewall in a VPC to block all domains except domains that are on an approved list. The company is concerned that if DNS Firewall is unresponsive, resources in the VPC might be affected if the network cannot resolve any DNS queries. To maintain application service level agreements, the company needs DNS queries to continue to resolve even if Route 53 Resolver does not receive a response from DNS Firewall.

Which change should a network engineer implement to meet these requirements?

- A. Update the DNS Firewall VPC configuration to disable fail open for the VPC.
- B. Update the DNS Firewall VPC configuration to enable fail open for the VPC.
- C. Create a new DHCP options set with parameter dns_firewall_fail_open=fals
- D. Associate the new DHCP options set with the VPC.
- E. Create a new DHCP options set with parameter dns_firewall_fail_open=tru
- F. Associate the new DHCP options set with the VPC.

Answer: B

NEW QUESTION 21

An AWS CloudFormation template is being used to create a VPC peering connection between two existing operational VPCs, each belonging to a different AWS account. All necessary components in the 'Remote' (receiving) account are already in place.

The template below creates the VPC peering connection in the Originating account. It contains these components:

AWSTemplateFormation Version: 2010-09-09 Parameters:

Originating VPCId: Type: String RemoteVPCId: Type: String

RemoteVPCAccountId: Type: String Resources:

newVPCPeeringConnection:

Type: 'AWS::EC2::VPCPeeringConnection'

Properties:

VpcdId: !Ref OriginatingVPCId PeerVpcId: !Ref RemoteVPCId PeerOwnerId: !Ref RemoteVPCAccountId

Which additional AWS CloudFormation components are necessary in the Originating account to create an operational cross-account VPC peering connection with AWS CloudFormation? (Select two.)

- A. Resources:NewEC2SecurityGroup:Type: AWS::EC2::SecurityGroup
- B. Resources:NetworkInterfaceToRemoteVPC:Type: "AWS::EC2NetworkInterface"
- C. Resources:newEC2Route:Type: AWS::EC2::Route
- D. Resources:VPCGatewayToRemoteVPC:Type: "AWS::EC2::VPCGatewayAttachment"
- E. Resources:newVPCPeeringConnection:Type: 'AWS::EC2VPCPeeringConnection'PeerRoleArn: !Ref PeerRoleArn

Answer: CE

Explanation:

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/AWS_EC2.html

NEW QUESTION 23

A company is deploying an application. The application is implemented in a series of containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use the Fargate launch type for its tasks. The containers will run workloads that require connectivity initiated over an SSL connection. Traffic must be able to flow to the application from other AWS accounts over private connectivity. The application must scale in a manageable way as more consumers use the application.

Which solution will meet these requirements?

- A. Choose a Gateway Load Balancer (GLB) as the type of load balancer for the ECS service
- B. Create a lifecycle hook to add new tasks to the target group from Amazon ECS as required to handle scaling
- C. Specify the GLB in the service definition
- D. Create a VPC peer for external AWS account
- E. Update the route tables so that the AWS accounts can reach the GLB.
- F. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service
- G. Create path-based routing rules to allow the application to target the containers that are registered in the target group
- H. Specify the ALB in the service definition
- I. Create a VPC endpoint service for the ALB. Share the VPC endpoint service with other AWS accounts.
- J. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS service
- K. Create path-based routing rules to allow the application to target the containers that are registered in the target group
- L. Specify the ALB in the service definition
- M. Create a VPC peer for the external AWS account
- N. Update the route tables so that the AWS accounts can reach the ALB.
- O. Choose a Network Load Balancer (NLB) as the type of load balancer for the ECS service
- P. Specify the NLB in the service definition
- Q. Create a VPC endpoint service for the NLB
- R. Share the VPC endpoint service with other AWS accounts.

Answer: D

NEW QUESTION 24

A company has created three VPCs: a production VPC, a nonproduction VPC, and a shared services VPC. The production VPC and the nonproduction VPC must each have communication with the shared services VPC. There must be no communication between the production VPC and the nonproduction VPC. A transit gateway is deployed to facilitate communication between VPCs.

Which route table configurations on the transit gateway will meet these requirements?

- A. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for only the shared services VPC
- B. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.
- C. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes for each VPC
- D. Create an additional route table with only the shared services VPC attachment associated with propagated routes from each VPC.
- E. Configure a route table with all the VPC attachments associated with propagated routes for only the shared services VPC. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.
- F. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes disabled
- G. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.

Answer: A

NEW QUESTION 27

A company is hosting an application on Amazon EC2 instances behind an Application Load Balancer. The instances are in an Amazon EC2 Auto Scaling group. Because of a recent change to a security group, external users cannot access the application.

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

Which solution will meet these requirements?

- A. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuration
- B. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.
- C. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration
- D. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.
- E. Configure Amazon GuardDuty to detect inconsistencies between the desired security group configuration and the current security group configuration
- F. Configure AWS OpsWorks for Chef to remediate noncompliant security groups.
- G. Configure an AWS Config rule to detect inconsistencies between the desired security group configuration and the current security group configuration
- H. Create an AWS Systems Manager Automation runbook to remediate noncompliant security groups.

Answer: D

Explanation:

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

NEW QUESTION 30

A company's development team has created a new product recommendation web service. The web service is hosted in a VPC with a CIDR block of 192.168.224.0/19. The company has deployed the web service on Amazon EC2 instances and has configured an Auto Scaling group as the target of a Network Load Balancer (NLB).

The company wants to perform testing to determine whether users who receive product recommendations spend more money than users who do not receive product recommendations. The company has a big sales event in 5 days and needs to integrate its existing production environment with the recommendation engine by then. The existing production environment is hosted in a VPC with a CIDR block of 192.168.128.0/17.

A network engineer must integrate the systems by designing a solution that results in the least possible disruption to the existing environments.

Which solution will meet these requirements?

- A. Create a VPC peering connection between the web service VPC and the existing production VPC
- B. Add a routing rule to the appropriate route table to allow data to flow to 192.168.224.0/19 from the existing production environment and to flow to

192.168.128.0/17 from the web service environmen

C. Configure the relevant security groups and ACLs to allow the systems to communicate.

D. Ask the development team of the web service to redeploy the web service into the production VPC and integrate the systems there.

E. Create a VPC endpoint service

F. Associate the VPC endpoint service with the NLB for the web service. Create an interface VPC endpoint for the web service in the existing production VPC.

G. Create a transit gateway in the existing production environmen

H. Create attachments to the production VPC and the web service VP

I. Configure appropriate routing rules in the transit gateway and VPC route tables for 192.168.224.0/19 and 192.168.128.0/17. Configure the relevant security groups and ACLs to allow the systems to communicate.

Answer: C

NEW QUESTION 31

A company hosts an application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company recently experienced a network security breach. A network engineer must collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application.

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

- A. Configure the ALB to store logs in an Amazon S3 bucket
- B. Download the files from Amazon S3, and use a spreadsheet application to analyze the logs.
- C. Configure the ALB to push logs to Amazon Kinesis Data Stream
- D. Use Amazon Kinesis Data Analytics to analyze the logs.
- E. Configure Amazon Kinesis Data Streams to stream data from the ALB to Amazon OpenSearch Service (Amazon Elasticsearch Service). Use search operations in Amazon OpenSearch Service (Amazon Elasticsearch Service) to analyze the data.
- F. Configure the ALB to store logs in an Amazon S3 bucket
- G. Use Amazon Athena to analyze the logs in Amazon S3.

Answer: D

Explanation:

The most operationally efficient solution to collect and analyze logs that include the client IP address, target IP address, target port, and user agent of each user that accesses the application would be to configure the ALB to store logs in an Amazon S3 bucket and use Amazon Athena to analyze the logs in Amazon S3 (Option D). This solution allows for quick and easy analysis of log data without requiring manual download or manipulation of log files.

NEW QUESTION 33

A company is planning to create a service that requires encryption in transit. The traffic must not be decrypted between the client and the backend of the service. The company will implement the service by using the gRPC protocol over TCP port 443. The service will scale up to thousands of simultaneous connections. The backend of the service will be hosted on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with the Kubernetes Cluster Autoscaler and the Horizontal Pod Autoscaler configured. The company needs to use mutual TLS for two-way authentication between the client and the backend.

Which solution will meet these requirements?

- A. Install the AWS Load Balancer Controller for Kubernetes
- B. Using that controller, configure a Network Load Balancer with a TCP listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
- C. Install the AWS Load Balancer Controller for Kubernetes
- D. Using that controller, configure an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the IP addresses of the backend service Pods.
- E. Create a target group
- F. Add the EKS managed node group's Auto Scaling group as a target. Create an Application Load Balancer with an HTTPS listener on port 443 to forward traffic to the target group.
- G. Create a target group
- H. Add the EKS managed node group's Auto Scaling group as a target
- I. Create a Network Load Balancer with a TLS listener on port 443 to forward traffic to the target group.

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html#target-gro>

NEW QUESTION 34

A software company offers a software-as-a-service (SaaS) accounting application that is hosted in the AWS Cloud. The application requires connectivity to the company's on-premises network. The company has two redundant 10 GB AWS Direct Connect connections between AWS and its on-premises network to accommodate the growing demand for the application.

The company already has encryption between its on-premises network and the colocation. The company needs to encrypt traffic between AWS and the edge routers in the colocation within the next few months. The company must maintain its current bandwidth.

What should a network engineer do to meet these requirements with the LEAST operational overhead?

- A. Deploy a new public VIF with encryption on the existing Direct Connect connection
- B. Reroute traffic through the new public VIF.
- C. Create a virtual private gateway. Deploy new AWS Site-to-Site VPN connections from on premises to the virtual private gateway. Reroute traffic from the Direct Connect private VIF to the new VPNs.
- D. Deploy a new pair of 10 GB Direct Connect connections with MACsec
- E. Configure MACsec on the edge router
- F. Reroute traffic to the new Direct Connect connection
- G. Decommission the original Direct Connect connections
- H. Deploy a new pair of 10 GB Direct Connect connections with MACsec
- I. Deploy a new public VIF on the new Direct Connect connection
- J. Deploy two AWS Site-to-Site VPN connections on top of the new public VIF
- K. Reroute traffic from the existing private VIF to the new Site-to-Site connection
- L. Decommission the original Direct Connect connections.

Answer: C

NEW QUESTION 37

An organization is using a VPC endpoint for Amazon S3. When the security group rules for a set of instances were initially configured, access was restricted to allow traffic only to the IP addresses of the Amazon S3 API endpoints in the region from the published JSON file. The application was working properly, but now is logging a growing number of timeouts when connecting with Amazon S3. No internet gateway is configured for the VPC.

Which solution will fix the connectivity failures with the LEAST amount of effort?

- A. Create a Lambda function to update the security group based on AmazonIPSpaceChanged notifications.
- B. Update the VPC routing to direct Amazon S3 prefix-list traffic to the VPC endpoint using the route table APIs.
- C. Update the application server's outbound security group to use the prefix-list for Amazon S3 in the same region.
- D. Create an additional VPC endpoint for Amazon S3 in the same route table to scale the concurrent connections to Amazon.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/subscribe-to-aws-public-ip-address-changes-via-amazon-sns/>

NEW QUESTION 40

A banking company is successfully operating its public mobile banking stack on AWS. The mobile banking stack is deployed in a VPC that includes private subnets and public subnets. The company is using IPv4 networking and has not deployed or supported IPv6 in the environment. The company has decided to adopt a third-party service provider's API and must integrate the API with the existing environment. The service provider's API requires the use of IPv6.

A network engineer must turn on IPv6 connectivity for the existing workload that is deployed in a private subnet. The company does not want to permit IPv6 traffic from the public internet and mandates that the company's servers must initiate all IPv6 connectivity. The network engineer turns on IPv6 in the VPC and in the private subnets.

Which solution will meet these requirements?

- A. Create an internet gateway and a NAT gateway in the VP
- B. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT gateway.
- C. Create an internet gateway and a NAT instance in the VP
- D. Add a route to the existing subnet route tables to point IPv6 traffic to the NAT instance.
- E. Create an egress-only Internet gateway in the VPAdd a route to the existing subnet route tables to point IPv6 traffic to the egress-only internet gateway.
- F. Create an egress-only internet gateway in the VP
- G. Configure a security group that denies all inbound traffic
- H. Associate the security group with the egress-only internet gateway.

Answer: C

NEW QUESTION 41

A company's network engineer is designing an active-passive connection to AWS from two on-premises data centers. The company has set up AWS Direct Connect connections between the on-premises data centers and AWS. From each location, the company is using a transit VIF that connects to a Direct Connect gateway that is associated with a transit gateway.

The network engineer must ensure that traffic from AWS to the data centers is routed first to the primary data center. The traffic should be routed to the failover data center only in the case of an outage.

Which solution will meet these requirements?

- A. Set the BGP community tag for all prefixes from the primary data center to 7224:7100. Set the BGP community tag for all prefixes from the failover data center to 7224:7300
- B. Set the BGP community tag for all prefixes from the primary data center to 7224:7300. Set the BGP community tag for all prefixes from the failover data center to 7224:7100
- C. Set the BGP community tag for all prefixes from the primary data center to 7224:9300. Set the BGP community tag for all prefixes from the failover data center to 7224:9100
- D. Set the BGP community tag for all prefixes from the primary data center to 7224:9100. Set the BGP community tag for all prefixes from the failover data center to 7224:9300

Answer: B

NEW QUESTION 42

A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin in an Amazon CloudFront distribution. The company wants to implement a custom authentication system that will provide a token for its authenticated customers.

The web application must ensure that the GET/POST requests come from authenticated customers before it delivers the content. A network engineer must design a solution that gives the web application the ability to identify authorized customers.

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

- A. Use the ALB to inspect the authorized token inside the GET/POST request payload
- B. Use an AWS Lambda function to insert a customized header to inform the web application of an authenticated customer request.
- C. Integrate AWS WAF with the ALB to inspect the authorized token inside the GET/POST request payload
- D. Configure the ALB listener to insert a customized header to inform the web application of an authenticated customer request.
- E. Use an AWS Lambda@Edge function to inspect the authorized token inside the GET/POST request payload
- F. Use the Lambda@Edge function also to insert a customized header to inform the web application of an authenticated customer request.
- G. Set up an EC2 instance that has a third-party packet inspection tool to inspect the authorized token inside the GET/POST request payload
- H. Configure the tool to insert a customized header to inform the web application of an authenticated customer request.

Answer: C

NEW QUESTION 44

An international company provides early warning about tsunamis. The company plans to use IoT devices to monitor sea waves around the world. The data that is collected by the IoT devices must reach the company's infrastructure on AWS as quickly as possible. The company is using three operation centers around the

world. Each operation center is connected to AWS through its own AWS Direct Connect connection. Each operation center is connected to the internet through at least two upstream internet service providers.

The company has its own provider-independent (PI) address space. The IoT devices use TCP protocols for reliable transmission of the data they collect. The IoT devices have both landline and mobile internet connectivity. The infrastructure and the solution will be deployed in multiple AWS Regions. The company will use Amazon Route 53 for DNS services.

A network engineer needs to design connectivity between the IoT devices and the services that run in the AWS Cloud.

Which solution will meet these requirements with the HIGHEST availability?

- A. Set up an Amazon CloudFront distribution with origin failover
- B. Create an origin group for each Region where the solution is deployed.
- C. Set up Route 53 latency-based routing
- D. Add latency alias record
- E. For the latency alias records, set the value of Evaluate Target Health to Yes.
- F. Set up an accelerator in AWS Global Accelerator
- G. Configure Regional endpoint groups and health checks.
- H. Set up Bring Your Own IP (BYOIP) addresses
- I. Use the same PI addresses for each Region where the solution is deployed.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/iot/automate-global-device-provisioning-with-aws-iot-core-and-amazon-route-53>

NEW QUESTION 49

A company's network engineer needs to design a new solution to help troubleshoot and detect network anomalies. The network engineer has configured Traffic Mirroring. However, the mirrored traffic is overwhelming the Amazon EC2 instance that is the traffic mirror target. The EC2 instance hosts tools that the company's security team uses to analyze the traffic. The network engineer needs to design a highly available solution that can scale to meet the demand of the mirrored traffic.

Which solution will meet these requirements?

- A. Deploy a Network Load Balancer (NLB) as the traffic mirror target
- B. Behind the NLB
- C. Deploy a fleet of EC2 instances in an Auto Scaling group
- D. Use Traffic Mirroring as necessary.
- E. Deploy an Application Load Balancer (ALB) as the traffic mirror target
- F. Behind the ALB, deploy a fleet of EC2 instances in an Auto Scaling group
- G. Use Traffic Mirroring only during non-business hours.
- H. Deploy a Gateway Load Balancer (GLB) as the traffic mirror target
- I. Behind the GLB
- J. Deploy a fleet of EC2 instances in an Auto Scaling group
- K. Use Traffic Mirroring as necessary.
- L. Deploy an Application Load Balancer (ALB) with an HTTPS listener as the traffic mirror target
- M. Behind the ALB
- N. Deploy a fleet of EC2 instances in an Auto Scaling group
- O. Use Traffic Mirroring only during active events or business hours.

Answer: A

NEW QUESTION 53

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

Which solution will meet these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a custom pattern to monitor the account for change
- B. Configure the rule to invoke an AWS Lambda function to identify noncompliant resource
- C. Update an Amazon DynamoDB table with the changes that are identified.
- D. Create custom metrics from Amazon CloudWatch log
- E. Use the metrics to invoke an AWS Lambda function to identify noncompliant resource
- F. Update an Amazon DynamoDB table with the changes that are identified.
- G. Record the current state of network resources by using AWS Config
- H. Create rules that reflect the desired configuration settings
- I. Set remediation for noncompliant resources.
- J. Record the current state of network resources by using AWS Systems Manager Inventory
- K. Use Systems Manager State Manager to enforce the desired configuration settings and to carry out remediation for noncompliant resources.

Answer: C

Explanation:

Recording the current state of network resources by using AWS Config would enable auditing and assessment of resource configurations and compliance.

Creating rules that reflect the desired configuration settings would enable evaluation of whether the network resources comply with network security policies.

Setting remediation for noncompliant resources would enable automatic correction of undesired configurations.

NEW QUESTION 58

A company operates its IT services through a multi-site hybrid infrastructure. The company deploys resources on AWS in the us-east-1 Region and in the eu-west-2 Region. The company also deploys resources in its own data centers that are located in the United States (US) and in the United Kingdom (UK). In both AWS Regions, the company uses a transit gateway to connect 15 VPCs to each other. The company has created a transit gateway peering connection between the two transit gateways. The VPC CIDR blocks do not overlap with each other or with IP addresses used within the data centers. The VPC CIDR prefixes can also be aggregated either on a Regional level or for the company's entire AWS environment.

The data centers are connected to each other by a private WAN connection. IP routing information is exchanged dynamically through Interior BGP (iBGP)

sessions. The data centers maintain connectivity to AWS through one AWS Direct Connect connection in the US and one Direct Connect connection in the UK. Each Direct Connect connection is terminated on a Direct Connect gateway and is associated with a local transit gateway through a transit VIF. Traffic follows the shortest geographical path from source to destination. For example, packets from the UK data center that are targeted to resources in eu-west-2 travel across the local Direct Connect connection. In cases of cross-Region data transfers, such as from the UK data center to VPCs in us-east-1, the private WAN connection must be used to minimize costs on AWS. A network engineer has configured each transit gateway association on the Direct Connect gateway to advertise VPC-specific CIDR IP prefixes only from the local Region. The routes toward the other Region must be learned through BGP from the routers in the other data center in the original, non-aggregated form.

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

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

- A. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connectio
- B. Add the company's entire AWS environment aggregate route to the list of subnets advertised through the local Direct Connect connection.
- C. Add the CIDR prefixes from the other Region VPCs and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connectio
- D. Configure data center routers to make routing decisions based on the BGP communities received.
- E. Add the aggregate IP prefix for the other Region and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection.
- F. Add the aggregate IP prefix for the company's entire AWS environment and the local VPC CIDR blocks to the list of subnets advertised through the local Direct Connect connection.
- G. Remove all the VPC CIDR prefixes from the list of subnets advertised through the local Direct Connect connectio
- H. Add both Regional aggregate IP prefixes to the list of subnets advertised through the Direct Connect connection on both sides of the networ
- I. Configure data center routers to make routing decisions based on the BGP communities received.

Answer: AD

NEW QUESTION 61

A company deploys a new web application on Amazon EC2 instances. The application runs in private subnets in three Availability Zones behind an Application Load Balancer (ALB). Security auditors require encryption of all connections. The company uses Amazon Route 53 for DNS and uses AWS Certificate Manager (ACM) to automate SSL/TLS certificate provisioning. SSL/TLS connections are terminated on the ALB.

The company tests the application with a single EC2 instance and does not observe any problems. However, after production deployment, users report that they can log in but that they cannot use the application. Every new web request restarts the login process.

What should a network engineer do to resolve this issue?

- A. Modify the ALB listener configuratio
- B. Edit the rule that forwards traffic to the target grou
- C. Change the rule to enable group-level stickines
- D. Set the duration to the maximum application session length.
- E. Replace the ALB with a Network Load Balance
- F. Create a TLS listene
- G. Create a new target group with the protocol type set to TLS Register the EC2 instance
- H. Modify the target group configuration by enabling the stickiness attribute.
- I. Modify the ALB target group configuration by enabling the stickiness attribut
- J. Use an application-based cooki
- K. Set the duration to the maximum application session length.
- L. Remove the AL
- M. Create an Amazon Route 53 rule with a failover routing policy for the application nam
- N. Configure ACM to issue certificates for each EC2 instance.

Answer: C

NEW QUESTION 62

A company has been using an outdated application layer protocol for communication among applications. The company decides not to use this protocol anymore and must migrate all applications to support a new protocol. The old protocol and the new protocol are TCP-based, but the protocols use different port numbers. After several months of work, the company has migrated dozens of applications that run on Amazon EC2 instances and in containers. The company believes that all the applications have been migrated, but the company wants to verify this belief. A network engineer needs to verify that no application is still using the old protocol.

Which solution will meet these requirements without causing any downtime?

- A. Use Amazon Inspector and its Network Reachability rules packag
- B. Wait until the analysis has finished running to find out which EC2 instances are still listening to the old port.
- C. Enable Amazon GuardDut
- D. Use the graphical visualizations to filter for traffic that uses the port of the old protoco
- E. Exclude all internet traffic to filter out occasions when the same port is used as an ephemeral port.
- F. Configure VPC flow logs to be delivered into an Amazon S3 bucke
- G. Use Amazon Athena to query the data and to filter for the port number that is used by the old protocol.
- H. Inspect all security groups that are assigned to the EC2 instances that host the application
- I. Remove the port of the old protocol if that port is in the list of allowed port
- J. Verify that the applications are operating properly after the port is removed from the security groups.

Answer: C

Explanation:

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

NEW QUESTION 65

A network engineer is designing a hybrid architecture that uses a 1 Gbps AWS Direct Connect connection between the company's data center and two AWS Regions: us-east-1 and eu-west-1. The VPCs in us-east-1 are connected by a transit gateway and need to access several on-premises databases. According to company policy, only one VPC in eu-west-1 can be connected to one on-premises server. The on-premises network segments the traffic between the databases

and the server.

How should the network engineer set up the Direct Connect connection to meet these requirements?

- A. Create one hosted connectio
- B. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use one Direc
- C. Connect gateway for both VIFs to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.
- D. Create one hosted connectio
- E. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use two Direct Connect gateways, one for each VIF, to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.
- F. Create one dedicated connectio
- G. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use one Direct Connect gateway for both VIFs to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.
- H. Create one dedicated connectio
- I. Use a transit VIF to connect to the transit gateway in us-east-1. Use a private VIF to connect to the VPC in eu-west-1. Use two Direct Connect gateways, one for each VIF, to route from the Direct Connect locations to the corresponding AWS Region along the path that has the lowest latency.

Answer: B

Explanation:

This solution meets the requirements of the company by using a single Direct Connect connection with two VIFs, one connected to the transit gateway in us-east-1 and the other connected to the VPC in eu-west-1. Two Direct Connect gateways are used, one for each VIF, to route traffic from the Direct Connect location to the corresponding AWS Region along the path that has the lowest latency. This setup ensures that traffic between the VPCs in us-east-1 and on-premises databases is routed through the transit gateway, while traffic between the VPC in eu-west-1 and the on-premises server is routed directly through the private VIF.

NEW QUESTION 67

A company uses a 4 Gbps AWS Direct Connect dedicated connection with a link aggregation group (LAG) bundle to connect to five VPCs that are deployed in the us-east-1 Region. Each VPC serves a different business unit and uses its own private VIF for connectivity to the on-premises environment. Users are reporting slowness when they access resources that are hosted on AWS.

A network engineer finds that there are sudden increases in throughput and that the Direct Connect connection becomes saturated at the same time for about an hour each business day. The company wants to know which business unit is causing the sudden increase in throughput. The network engineer must find out this information and implement a solution to resolve the problem.

Which solution will meet these requirements?

- A. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observe
- B. Create a new 10 Gbps dedicated connectio
- C. Shift traffic from the existing dedicated connection to the new dedicated connection.
- D. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observe
- E. Upgrade the bandwidth of the existing dedicated connection to 10 Gbps.
- F. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observe
- G. Upgrade the existing dedicated connection to a 5 Gbps hosted connection.
- H. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observed.Create a new 10 Gbps dedicated connectio
- I. Shift traffic from the existing dedicated connection to the new dedicated connection.

Answer: A

Explanation:

To meet the requirements of finding out which business unit is causing the sudden increase in throughput and resolving the problem, the network engineer should review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observed (Option B). After identifying the VIF that is causing the issue, they can upgrade the bandwidth of the existing dedicated connection to 10 Gbps to resolve the problem (Option B).

NEW QUESTION 72

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

Which solution will meet these requirements MOST cost-effectively?

- A. Use an Application Load Balancer (ALB)-type target group for a Network Load Balancer (NLB). Create an EC2 Auto Scaling grou
- B. Attach the Auto Scaling group to the AL
- C. Set up the IoT devices to connect to the IP addresses of the NLB.
- D. Use an AWS Global Accelerator accelerator with an Application Load Balancer (ALB) endpoint
- E. Create an EC2 Auto Scaling grou
- F. Attach the Auto Scaling group to the ALSet up the IoT devices to connect to the IP addresses of the accelerator.
- G. Use a Network Load Balancer (NLB). Create an EC2 Auto Scaling grou
- H. Attach the Auto Scaling group to the NL
- I. Set up the IoT devices to connect to the IP addresses of the NLB.
- J. Use an AWS Global Accelerator accelerator with a Network Load Balancer (NLB) endpoint
- K. Create anEC2 Auto Scaling grou
- L. Attach the Auto Scaling group to the NL
- M. 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.

NEW QUESTION 74

A company wants to improve visibility into its AWS environment. The AWS environment consists of multiple VPCs that are connected to a transit gateway. The transit gateway connects to an on-premises data center through an AWS Direct Connect gateway and a pair of redundant Direct Connect connections that use transit VIFs. The company must receive notification each time a new route is advertised to AWS from on premises over Direct Connect. What should a network engineer do to meet these requirements?

- A. Enable Amazon CloudWatch metrics on Direct Connect to track the received route
- B. Configure a CloudWatch alarm to send notifications when routes change.
- C. Onboard Transit Gateway Network Manager to Amazon CloudWatch Logs Insight
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to send notifications when routes change.
- E. Configure an AWS Lambda function to periodically check the routes on the Direct Connect gateway and to send notifications when routes change.
- F. Enable Amazon CloudWatch Logs on the transit VIFs to track the received route
- G. Create a metric filter Set an alarm on the filter to send notifications when routes change.

Answer: B

Explanation:

<https://docs.aws.amazon.com/network-manager/latest/cloudwan/cloudwan-cloudwatch-events.html>

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

NEW QUESTION 77

An organization is replacing a tape backup system with a storage gateway. there is currently no connectivity to AWS. Initial testing is needed. What connection option should the organization use to get up and running at minimal cost?

- A. Use an internet connection.
- B. Set up an AWS VPN connection.
- C. Provision an AWS Direct Connection private virtual interface.
- D. Provision a Direct Connect public virtual interface.

Answer: A

NEW QUESTION 81

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

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

What should a network engineer do to meet these requirements?

- A. Create a new Route 53 Resolver inbound endpoint in the shared services VP
- B. Create forwarding rules for the on-premises hosted domain
- C. Associate the rules with the new Resolver endpoint and each application VP
- D. Update each application VPC's DHCP configuration to point DNS resolution to the new Resolver endpoint.
- E. Create a new Route 53 Resolver outbound endpoint in the shared services VP
- F. Create forwarding rules for the on-premises hosted domain
- G. Associate the rules with the new Resolver endpoint and each application VPC.
- H. Create a new Route 53 Resolver outbound endpoint in the shared services VPC
- I. Associate the rules with the new Resolver endpoint and each application VP
- J. Create a new Route 53 Resolver inbound endpoint in the shared services VP
- K. Create forwarding rules for the on-premises hosted domain
- L. Associate the rules with the new Resolver endpoint and each application VPC.

Answer: B

Explanation:

Creating a new Route 53 Resolver outbound endpoint in the shared services VPC would enable forwarding of DNS queries from the VPC to on-premises¹.

Creating forwarding rules for the on-premises hosted domains would enable specifying which domain names are forwarded to the on-premises DNS servers².

Associating the rules with the new Resolver endpoint and each application VPC would enable applying the rules to the VPCs². This solution would not affect the default DNS resolution behavior of Route 53 Resolver for local VPC domain names and domains that are hosted in Route 53 private hosted zones³.

NEW QUESTION 82

A company has several production applications across different accounts in the AWS Cloud. The company operates from the us-east-1 Region only. Only certain partner companies can access the applications. The applications are running on Amazon EC2 instances that are in an Auto Scaling group behind an Application Load Balancer (ALB). The EC2 instances are in private subnets and allow traffic only from the ALB. The ALB is in a public subnet and allows inbound traffic only from partner network IP address ranges over port 80.

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

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

- A. Create an Amazon DynamoDB table to maintain all IP address ranges and security groups that need to be update
- B. Update the DynamoDB table with the new IP address range when the company adds a new partne
- C. Invoke an AWS Lambda function to read new IP address ranges and security groups from the DynamoDB table to update the security group
- D. Deploy this solution in all accounts.
- E. Create a new prefix lis
- F. Add all allowed IP address ranges to the prefix lis
- G. Use Amazon EventBridge (Amazon CloudWatch Events) rules to invoke an AWS Lambda function to update security groups whenever a new IP address range

- is added to the prefix list
- H. Deploy this solution in all accounts.
- I. Create a new prefix list
- J. Add all allowed IP address ranges to the prefix list
- K. Share the prefix list across different accounts by using AWS Resource Access Manager (AWS RAM). Update security groups to use the prefix list instead of the partner IP address range
- L. Update the prefix list with the new IP address range when the company adds a new partner.
- M. Create an Amazon S3 bucket to maintain all IP address ranges and security groups that need to be update
- N. Update the S3 bucket with the new IP address range when the company adds a new partner
- O. Invoke an AWS Lambda function to read new IP address ranges and security groups from the S3 bucket to update the security group
- P. Deploy this solution in all accounts.

Answer: C

Explanation:

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

NEW QUESTION 86

A security team is performing an audit of a company's AWS deployment. The security team is concerned that two applications might be accessing resources that should be blocked by network ACLs and security groups. The applications are deployed across two Amazon Elastic Kubernetes Service (Amazon EKS) clusters that use the Amazon VPC Container Network Interface (CNI) plugin for Kubernetes. The clusters are in separate subnets within the same VPC and have a Cluster Autoscaler configured.

The security team needs to determine which POD IP addresses are communicating with which services throughout the VPC. The security team wants to limit the number of flow logs and wants to examine the traffic from only the two applications.

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

- A. Create VPC flow logs in the default format
- B. Create a filter to gather flow logs only from the EKS nodes. Include the srcaddr field and the dstaddr field in the flow logs.
- C. Create VPC flow logs in a custom format
- D. Set the EKS nodes as the resource. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.
- E. Create VPC flow logs in a custom format
- F. Set the application subnets as resource
- G. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.
- H. Create VPC flow logs in a custom format
- I. Create a filter to gather flow logs only from the EKS nodes. Include the pkt-srcaddr field and the pkt-dstaddr field in the flow logs.

Answer: D

NEW QUESTION 90

A company delivers applications over the internet. An Amazon Route 53 public hosted zone is the authoritative DNS service for the company and its internet applications, all of which are offered from the same domain name.

A network engineer is working on a new version of one of the applications. All the application's components are hosted in the AWS Cloud. The application has a three-tier design. The front end is delivered through Amazon EC2 instances that are deployed in public subnets with Elastic IP addresses assigned. The backend components are deployed in private subnets from RFC1918.

Components of the application need to be able to access other components of the application within the application's VPC by using the same host names as the host names that are used over the public internet. The network engineer also needs to accommodate future DNS changes, such as the introduction of new host names or the retirement of DNS entries.

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

- A. Add a geoproximity routing policy in Route 53.
- B. Create a Route 53 private hosted zone for the same domain name. Associate the application's VPC with the new private hosted zone.
- C. Enable DNS hostnames for the application's VPC.
- D. Create entries in the private hosted zone for each name in the public hosted zone by using the corresponding private IP addresses.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs when AWS CloudTrail logs a Route 53 API call to the public hosted zone.
- F. Create an AWS Lambda function as the target of the rule.
- G. Configure the function to use the event information to update the private hosted zone.
- H. Add the private IP addresses in the existing Route 53 public hosted zone.

Answer: BCD

NEW QUESTION 92

A company uses AWS Direct Connect to connect its corporate network to multiple VPCs in the same AWS account and the same AWS Region. Each VPC uses its own private VIF and its own virtual LAN on the Direct Connect connection. The company has grown and will soon surpass the limit of VPCs and private VIFs for each connection.

What is the MOST scalable way to add VPCs with on-premises connectivity?

- A. Provision a new Direct Connect connection to handle the additional VPC
- B. Use the new connection to connect additional VPCs.
- C. Create virtual private gateways for each VPC that is over the service quota
- D. Use AWS Site-to-Site VPN to connect the virtual private gateways to the corporate network.
- E. Create a Direct Connect gateway, and add virtual private gateway associations to the VPC
- F. Configure a private VIF to connect to the corporate network.
- G. Create a transit gateway, and attach the VPC
- H. Create a Direct Connect gateway, and associate it with the transit gateway
- I. Create a transit VIF to the Direct Connect gateway.

Answer: D

Explanation:

When a company requires connectivity to multiple VPCs over AWS Direct Connect, a scalable solution is to use a transit gateway. A transit gateway is a hub that can interconnect multiple VPCs and VPN connections. The VPCs can communicate with each other over the transit gateway, and on-premises networks can communicate with the VPCs through the Direct Connect gateway. This solution provides a central point of management and simplifies the configuration of network routing. By associating the Direct Connect gateway with the transit gateway, traffic between the VPCs and the on-premises network can be routed through the Direct Connect connection.

NEW QUESTION 96

A company recently migrated its Amazon EC2 instances to VPC private subnets to satisfy a security compliance requirement. The EC2 instances now use a NAT gateway for internet access. After the migration, some long-running database queries from private EC2 instances to a publicly accessible third-party database no longer receive responses. The database query logs reveal that the queries successfully completed after 7 minutes but that the client EC2 instances never received the response.

Which configuration change should a network engineer implement to resolve this issue?

- A. Configure the NAT gateway timeout to allow connections for up to 600 seconds.
- B. Enable enhanced networking on the client EC2 instances.
- C. Enable TCP keepalive on the client EC2 instances with a value of less than 300 seconds.
- D. Close idle TCP connections through the NAT gateway.

Answer: C

Explanation:

When a TCP connection is idle for a long time, it may be terminated by network devices, including the NAT gateway. By enabling TCP keepalive, the client EC2 instances can periodically send packets to the third-party database to indicate that the connection is still active, preventing it from being terminated prematurely.

NEW QUESTION 100

An Australian ecommerce company hosts all of its services in the AWS Cloud and wants to expand its customer base to the United States (US). The company is targeting the western US for the expansion.

The company's existing AWS architecture consists of four AWS accounts with multiple VPCs deployed in the ap-southeast-2 Region. All VPCs are attached to a transit gateway in ap-southeast-2. There are dedicated VPCs for each application service. The company also has VPCs for centralized security features such as proxies, firewalls, and logging.

The company plans to duplicate the infrastructure from ap-southeast-2 to the us-west-1 Region. A network engineer must establish connectivity between the various applications in the two Regions. The solution must maximize bandwidth, minimize latency and minimize operational overhead.

Which solution will meet these requirements?

- A. Create VPN attachments between the two transit gateway
- B. Configure the VPN attachments to use BGP routing between the two transit gateways.
- C. Peer the transit gateways in each Region
- D. Configure routing between the two transit gateways for each Region's IP addresses.
- E. Create a VPN server in a VPC in each Region
- F. Update the routing to point to the VPN servers for the IP addresses in alternate Regions.
- G. Attach the VPCs in us-west-1 to the transit gateway in ap-southeast-2.

Answer: B

Explanation:

Peering the transit gateways in each region would establish a private network connection between the two regions, allowing the company to route traffic between the VPCs in different regions without going over the public internet. This would help minimize latency and maximize bandwidth while reducing the operational overhead of managing multiple VPN connections.

NEW QUESTION 102

A retail company is running its service on AWS. The company's architecture includes Application Load Balancers (ALBs) in public subnets. The ALB target groups are configured to send traffic to backend Amazon EC2 instances in private subnets. These backend EC2 instances can call externally hosted services over the internet by using a NAT gateway.

The company has noticed in its billing that NAT gateway usage has increased significantly. A network engineer needs to find out the source of this increased usage.

Which options can the network engineer use to investigate the traffic through the NAT gateway? (Choose two.)

- A. Enable VPC flow logs on the NAT gateway's elastic network interface
- B. Publish the logs to a log group in Amazon CloudWatch Log
- C. Use CloudWatch Logs Insights to query and analyze the logs.
- D. Enable NAT gateway access log
- E. Publish the logs to a log group in Amazon CloudWatch Log
- F. Use CloudWatch Logs Insights to query and analyze the logs.
- G. Configure Traffic Mirroring on the NAT gateway's elastic network interface
- H. Send the traffic to an additional EC2 instance
- I. Use tools such as tcpdump and Wireshark to query and analyze the mirrored traffic.
- J. Enable VPC flow logs on the NAT gateway's elastic network interface
- K. Publish the logs to an Amazon S3 bucket
- L. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure
- M. Use Athena to query and analyze the logs.
- N. Enable NAT gateway access log
- O. Publish the logs to an Amazon S3 bucket
- P. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure
- Q. Use Athena to query and analyze the logs.

Answer: AD

Explanation:

To investigate the increased usage of a NAT gateway in a VPC architecture with ALBs and backend EC2 instances, a network engineer can use the following options:

➤ Enable VPC flow logs on the NAT gateway's elastic network interface and publish the logs to a log group in Amazon CloudWatch Logs. Use CloudWatch Logs Insights to query and analyze the logs.

(Option A)

➤ Enable VPC flow logs on the NAT gateway's elastic network interface and publish the logs to an Amazon S3 bucket. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure and use Athena to query and analyze the logs. (Option D)

These options allow for detailed analysis of traffic through the NAT gateway to identify the source of increased usage.

NEW QUESTION 107

A network engineer has deployed an Amazon EC2 instance in a private subnet in a VPC. The VPC has no public subnet. The EC2 instance hosts application code that sends messages to an Amazon Simple Queue Service (Amazon SQS) queue. The subnet has the default network ACL with no modification applied. The EC2 instance has the default security group with no modification applied.

The SQS queue is not receiving messages.

Which of the following are possible causes of this problem? (Choose two.)

- A. The EC2 instance is not attached to an IAM role that allows write operations to Amazon SQS.
- B. The security group is blocking traffic to the IP address range used by Amazon SQS
- C. There is no interface VPC endpoint configured for Amazon SQS
- D. The network ACL is blocking return traffic from Amazon SQS
- E. There is no route configured in the subnet route table for the IP address range used by Amazon SQS

Answer: CE

NEW QUESTION 109

A government contractor is designing a multi-account environment with multiple VPCs for a customer. A network security policy requires all traffic between any two VPCs to be transparently inspected by a third-party appliance.

The customer wants a solution that features AWS Transit Gateway. The setup must be highly available across multiple Availability Zones, and the solution needs to support automated failover. Furthermore, asymmetric routing is not supported by the inspection appliances.

Which combination of steps is part of a solution that meets these requirements? (Choose two.)

- A. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VP
- B. Connect the inspection VPC to the transit gateway by using a VPCattachmen
- C. Create a target group, and register the appliances with the target grou
- D. Create a Network Load Balancer (NLB), and set it up to forward to the newly created target grou
- E. Configure a default route in the inspection VPCs transit gateway subnet toward the NLB.
- F. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VP
- G. Connect the inspection VPC to the transit gateway by using a VPC attachmen
- H. Create a target group, and register the appliances with the target grou
- I. Create a Gateway Load Balancer, and set it up to forward to the newly created target grou
- J. Configure a default route in the inspection VPC's transit gateway subnet toward the Gateway Load Balancer endpoint.
- K. Configure two route tables on the transit gatewa
- L. Associate one route table with all the attachments of the application VPC
- M. Associate the other route table with the inspection VPC's attachmen
- N. Propagate all VPC attachments into the inspection route tabl
- O. Define a static default route in the application route tabl
- P. Enable appliance mode on the attachment that connects the inspection VPC.
- Q. Configure two route tables on the transit gatewa
- R. Associate one route table with all the attachments of the application VPC
- S. Associate the other route table with the inspection VPCs attachmen
- T. Propagate all VPC attachments into the application route tabl
- . Define a static default route in the inspection route tabl
- . Enable appliance mode on the attachment that connects the inspection VPC.
- . Configure one route table on the transit gatewa
- . Associate the route table with all the VPC
- . Propagate all VPC attachments into the route tabl
- . Define a static default route in the route table.

Answer: BC

NEW QUESTION 110

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

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

Which solution will meet these requirements?

- A. Create a Network Load Balance
- B. Create a target grou
- C. Set the protocol to TCP and the port to 443 for the target grou
- D. Turn on session affinity (sticky sessions). Register the EC2 instances as target
- E. Create a listene
- F. Set the protocol to TCP and the port to 443 for the listene
- G. Deploy SSL certificates to the EC2 instances.
- H. Create an Application Load Balance

- I. Create a target grou
- J. Set the protocol to HTTP and the port to 80 for the target grou
- K. Turn on session affinity (sticky sessions) with an application-based cookie polic
- L. Register the EC2 instances as target
- M. Create an HTTPS listene
- N. Set the default action to forward to the target grou
- O. Use AWS Certificate Manager (ACM) to create a certificatefor the listener.
- P. Create a Network Load Balance
- Q. Create a target grou
- R. Set the protocol to TLS and the port to 443 for the target grou
- S. Turn on session affinity (sticky sessions). Register the EC2 instances as target
- T. Create a listene
 - . Set the protocol to TLS and the port to 443 for the listene
 - . Use AWS Certificate Manager (ACM) to create a certificate for the application.
 - . Create an Application Load Balance
 - . Create a target grou
 - . Set the protocol to HTTPS and the port to 443 for the target grou
 - . Turn on session affinity (sticky sessions) with an application-based cookie polic
 - . Register the EC2 instances as target
 - . Create an HTTP listene
 - . Set the port to 443 for the listene
 - . Set the default action to forward to the target group.

Answer: A

NEW QUESTION 111

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your ANS-C01 Exam with Our Prep Materials Via below:

<https://www.certleader.com/ANS-C01-dumps.html>