ANS-C01 Q&As

AWS Certified Advanced Networking Specialty Exam

Pass Amazon ANS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/ans-c01.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A company has deployed a software-defined WAN (SD-WAN) solution to interconnect all of its offices. The company is migrating workloads toAWS and needs to extend its SD-WAN solution to support connectivity to these workloads.A network engineer plans to deploy AWS Transit Gateway Connect and two SD-WAN virtual appliances to provide this connectivity. Accordingto company policies, only a single SD-WAN virtual appliance can handle traffic from AWS workloads at a given time.How should the network engineer configure routing to meet these requirements?

A. Add a static default route in the transit gateway route table to point to the secondary SD-WAN virtual appliance. Add routes that aremore specific to point to the primary SD-WAN virtual appliance.

B. Configure the BGP community tag 7224:7300 on the primary SD-WAN virtual appliance for BGP routes toward the transit gateway.

C. Configure the AS_PATH prepend attribute on the secondary SD-WAN virtual appliance for BGP routes toward the transit gateway.

D. Disable equal-cost multi-path (ECMP) routing on the transit gateway for Transit Gateway Connect.

Correct Answer: C

A - incorrect , static routes are not possible in TGW

B- incorrect, these BGP communities are used for BGP over DX

C- correct , AS_PATH prepending is a standard BGP way of influencing return traffic for advertised prefixes and SDWAN supports this.

D- incorrect, disabling ECMP will make sure the SDWAN>TGW traffic is not load shared, but the return traffic TGW>SDWAN is not affected and therefore both appliances will process traffic.

QUESTION 2

A company runs an application on Amazon EC2 instances. A network engineer implements a NAT gateway in the application\\'s VPC to replaceself-managed NAT instances. After the network engineer shifts traffic from the self-managed NAT instances to the NAT gateway, users begin toreport issues.During troubleshooting, the network engineer discovers that the connection to the application is closing after approximately 6 minutes ofinactivity.What should the network engineer do to resolve this issue?

A. Check for increases in the IdleTimeoutCount Amazon CloudWatch metric for the NAT gateway. Configure TCP keepalive on theapplication EC2 instances.

B. Check for increases in the ErrorPortAllocation Amazon CloudWatch metric for the NAT gateway. Configure an HTTP timeout value on theapplication EC2 instances.

C. Check for increases in the PacketsDropCount Amazon CloudWatch metric for the NAT gateway. Configure an HTTPS timeout value onthe application EC2 instances.

D. Check for decreases in the ActiveConnectionCount Amazon CloudWatch metric for the NAT gateway. Configure UDP keepalive on theapplication EC2 instances.

Correct Answer: A

Internet connection drops after 350 seconds

Problem

Your instances can access the internet, but the connection drops after 350 seconds.

Cause

If a connection that\\'s using a NAT gateway is idle for 350 seconds or more, the connection times out.

When a connection times out, a NAT gateway returns an RST packet to any resources behind the NAT gateway that attempt to continue the connection (it does not send a FIN packet).

Solution

To prevent the connection from being dropped, you can initiate more traffic over the connection. Alternatively, you can enable TCP keepalive on the instance with a value less than 350 seconds.

QUESTION 3

A company is migrating its containerized application to AWS. For the architecture the company will have an ingress VPC with a Network LoadBalancer (NLB) to distribute the traffic to front-end pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The front end of theapplication will determine which user is requesting access and will send traffic to 1 of 10 services VPCs. Each services VPC will include anNLB that distributes traffic to the services pods in an EKS cluster.The company is concerned about overall cost. User traffic will be responsible for more than 10 TB of data transfer from the ingress VPC toservices VPCs every month. A network engineer needs to recommend how to design the communication between the VPCs.Which solution will meet these requirements at the LOWEST cost?

A. Create a transit gateway. Peer each VPC to the transit gateway. Use zonal DNS names for the NLB in the services VPCs to minimizecross-AZ traffic from the ingress VPC to the services VPCs.

B. Create an AWS PrivateLink endpoint in every Availability Zone in the ingress VPC. Each PrivateLink endpoint will point to the zonal DNSentry of the NLB in the services VPCs.

C. Create a VPC peering connection between the ingress VPC and each of the 10 services VPCs. Use zonal DNS names for the NLB in theservices VPCs to minimize cross-AZ traffic from the ingress VPC to the services VPCs.

D. Create a transit gateway. Peer each VPC to the transit gateway. Turn off cross-AZ load balancing on the transit gateway. Use RegionalDNS names for the NLB in the services VPCs.

Correct Answer: C

VPC peering offers the lowest overall cost when compared to other options for inter-VPC connectivity. https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/vpc-to-vpc-connectivity.html

There is no such thing as "TG peering"; there are VPC peering and TG attachments.

QUESTION 4

A company has deployed a critical application on a fleet of Amazon EC2 instances behind an Application Load Balancer. The application mustalways be reachable on port 443 from the public internet. The application recently had an outage that resulted from an incorrect change tothe EC2 security group.A network engineer needs to automate a way to verify the network connectivity between the public internet and the EC2 instances whenever achange is made to the

security group. The solution also must notify the network engineer when the change affects the connection.Which solution will meet these requirements?

A. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture REJECT traffic on port 443. Publish the flow logrecords to a log group in Amazon CloudWatch Logs. Create a CloudWatch Logs metric filter for the log group for rejected traffic. Create analarm to notify the network engineer.

B. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture all traffic on port 443. Publish the flow logrecords to a log group in Amazon CloudWatch Logs. Create a CloudWatch Logs metric filter for the log group for all traffic. Create an alarmto notify the network engineer

C. Create a VPC Reachability Analyzer path on port 443. Specify the security group as the source. Specify the EC2 instances as thedestination. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to thesecurity group affects the connection. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNStopic in case the analyses fail Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when achange to the security group occurs.

D. Create a VPC Reachability Analyzer path on port 443. Specify the internet gateway of the VPC as the source. Specify the EC2 instancesas the destination. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change tothe security group affects the connection. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to theSNS topic in case the analyses fail. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function whena change to the security group occurs.

Correct Answer: D

https://aws.amazon.com/blogs/networking-and-content-delivery/automating-connectivity-assessments-with-vpc-reachability-analyzer/

QUESTION 5

A global company runs business applications in the us-east-1 Region inside a VPC. One of the company\\'s regional offices in London uses avirtual private gateway for an AWS Site-to-Site VPN connection tom the VPC. The company has configured a transit gateway and has set uppeering between the VPC and other VPCs that various departments in the company use.Employees at the London office are experiencing latency issues when they connect to the business applications.What should a network engineer do to reduce this latency?

A. Create a new Site-to-Site VPN connection. Set the transit gateway as the target gateway. Enable acceleration on the new Site-to-SiteVPN connection. Update the VPN device in the London office with the new connection details.

B. Modify the existing Site-to-Site VPN connection by setting the transit gateway as the target gateway. Enable acceleration on theexisting Site-to-Site VPN connection.

C. Create a new transit gateway in the eu-west-2 (London) Region. Peer the new transit gateway with the existing transit gateway. Modifythe existing Site-to-Site VPN connection by setting the new transit gateway as the target gateway.

D. Create a new AWS Global Accelerator standard accelerator that has an endpoint of the Site-to-Site VPN connection. Update the VPNdevice in the London office with the new connection details.

Correct Answer: A

https://docs.aws.amazon.com/vpn/latest/s2svpn/accelerated-vpn.html

ANS-C01 Practice Test ANS-C01 Exam Questions ANS-C01 Braindumps