



**Amazon**

## **Exam Questions AWS-Certified-Advanced-Networking-Specialty**

Amazon AWS Certified Advanced Networking - Specialty

#### NEW QUESTION 1

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 2

A company is running multiple workloads on Amazon EC2 instances in public subnets. In a recent incident, an attacker exploited an application vulnerability on one of the EC2 instances to gain access to the instance. The company fixed the application and launched a replacement EC2 instance that contains the updated application.

The attacker used the compromised application to spread malware over the internet. The company became aware of the compromise through a notification from AWS. The company needs the ability to identify when an application that is deployed on an EC2 instance is spreading malware.

Which solution will meet this requirement with the LEAST operational effort?

- A. Use Amazon GuardDuty to analyze traffic patterns by inspecting DNS requests and VPC flow logs.
- B. Use Amazon GuardDuty to deploy AWS managed decoy systems that are equipped with the most recent malware signatures.
- C. Set up a Gateway Load Balance
- D. Run an intrusion detection system (IDS) appliance from AWS Marketplace on Amazon EC2 for traffic inspection.
- E. Configure Amazon Inspector to perform deep packet inspection of outgoing traffic.

**Answer:** A

#### Explanation:

This solution involves using Amazon GuardDuty to monitor network traffic and analyze DNS requests and VPC flow logs for suspicious activity. This will allow the company to identify when an application is spreading malware by monitoring the network traffic patterns associated with the instance. GuardDuty is a fully managed threat detection service that continuously monitors for malicious activity and unauthorized behavior in your AWS accounts and workloads. It requires minimal setup and configuration and can be integrated with other AWS services for automated remediation. This solution requires the least operational effort compared to the other options

#### NEW QUESTION 3

An ecommerce company is hosting a web application on Amazon EC2 instances to handle continuously changing customer demand. The EC2 instances are part of an Auto Scaling group. The company wants to implement a solution to distribute traffic from customers to the EC2 instances. The company must encrypt all traffic at all stages between the customers and the application servers. No decryption at intermediate points is allowed.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB). Add an HTTPS listener to the AL
- B. Configure the Auto Scaling group to register instances with the ALB's target group.
- C. Create an Amazon CloudFront distributio
- D. Configure the distribution with a custom SSL/TLS certificat
- E. Set the Auto Scaling group as the distribution's origin.
- F. Create a Network Load Balancer (NLB). Add a TCP listener to the NL
- G. Configure the Auto Scaling group to register instances with the NLB's target group.
- H. Create a Gateway Load Balancer (GLB). Configure the Auto Scaling group to register instances with the GLB's target group.

**Answer:** C

#### Explanation:

To distribute traffic from customers to EC2 instances in an Auto Scaling group and encrypt all traffic at all stages between the customers and the application servers without decryption at intermediate points, the company should create a Network Load Balancer (NLB) with a TCP listener and configure the Auto Scaling group to register instances with the NLB's target group (Option C). This solution allows for end-to-end encryption of traffic without decryption at intermediate points.

#### NEW QUESTION 4

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

**Answer:** A

#### NEW QUESTION 5

A network engineer is designing the architecture for a healthcare company's workload that is moving to the AWS Cloud. All data to and from the on-premises environment must be encrypted in transit. All traffic also must be inspected in the cloud before the traffic is allowed to leave the cloud and travel to the on-premises environment or to the internet.

The company will expose components of the workload to the internet so that patients can reserve appointments. The architecture must secure these components and protect them against DDoS attacks. The architecture also must provide protection against financial liability for services that scale out during a DDoS event. Which combination of steps should the network engineer take to meet all these requirements for the workload? (Choose three.)

- A. Use Traffic Mirroring to copy all traffic to a fleet of traffic capture appliances.
- B. Set up AWS WAF on all network components.
- C. Configure an AWS Lambda function to create Deny rules in security groups to block malicious IP addresses.
- D. Use AWS Direct Connect with MACsec support for connectivity to the cloud.
- E. Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection.
- F. Configure AWS Shield Advanced and ensure that it is configured on all public assets.

**Answer:** DEF

#### Explanation:

To meet the requirements for the healthcare company's workload that is moving to the AWS Cloud, the network engineer should take the following steps:

- Use AWS Direct Connect with MACsec support for connectivity to the cloud to ensure that all data to and from the on-premises environment is encrypted in transit (Option D).
  - Use Gateway Load Balancers to insert third-party firewalls for inline traffic inspection to inspect all traffic in the cloud before it is allowed to leave (Option E).
  - Configure AWS Shield Advanced and ensure that it is configured on all public assets to secure components exposed to the internet against DDoS attacks and provide protection against financial liability for services that scale out during a DDoS event (Option F).
- These steps will help ensure that all data is encrypted in transit, all traffic is inspected before leaving the cloud, and components exposed to the internet are secured against DDoS attacks.

#### NEW QUESTION 6

A company has its production VPC (VPC-A) in the eu-west-1 Region in Account 1. VPC-A is attached to a transit gateway (TGW-A) that is connected to an on-premises data center in Dublin, Ireland, by an AWS Direct Connect transit VIF that is configured for an AWS Direct Connect gateway. The company also has a staging VPC (VPC-B) that is attached to another transit gateway (TGW-B) in the eu-west-2 Region in Account 2.

A network engineer must implement connectivity between VPC-B and the on-premises data center in Dublin. Which solutions will meet these requirements? (Choose two.)

- A. Configure inter-Region VPC peering between VPC-A and VPC-B.
- B. Add the required VPC peering route
- C. Add the VPC-B CIDR block in the allowed prefixes on the Direct Connect gateway association.
- D. Associate TGW-B with the Direct Connect gateway
- E. Advertise the VPC-B CIDR block under the allowed prefixes.
- F. Configure another transit VIF on the Direct Connect connection and associate TGW-B.
- G. Advertise the VPC-B CIDR block under the allowed prefixes.
- H. Configure inter-Region transit gateway peering between TGW-A and TGW-B.
- I. Add the peering routes in the transit gateway route table
- J. Add both the VPC-A and the VPC-B CIDR block under the allowed prefix list in the Direct Connect gateway association.
- K. Configure an AWS Site-to-Site VPN connection over the transit VIF to TGW-B as a VPN attachment.

**Answer:** BC

#### Explanation:

\* B. Associate TGW-B with the Direct Connect gateway. Advertise the VPC-B CIDR block under the allowed prefixes. This will allow traffic from VPC-B to be sent over the Direct Connect connection to the on-premises data center via TGW-B. C. Configure another transit VIF on the Direct Connect connection and associate TGW-B. Advertise the VPC-B CIDR block under the allowed prefixes. This will enable the use of the Direct Connect connection for VPC-B's traffic by connecting TGW-B to the Direct Connect gateway.

#### NEW QUESTION 7

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 Gateway
- 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 VPC
- E. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access
- F. Ensure that private DNS is turned off
- G. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway
- 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 VPC. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access
- L. Ensure that private DNS is turned off
- M. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway
- N. Create an Amazon Route 53 private hosted zone with a full service endpoint name for Amazon S3 and Systems Manager
- 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 VPC
- R. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to access
- S. Connect all the VPCs to the central shared services VPC by using AWS Transit Gateway
- 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 connection. Interface VPC endpoints are powered by AWS PrivateLink, a technology that enables private access to AWS services. Amazon S3 and AWS Systems Manager support interface VPC endpoints. By turning off private DNS, the interface VPC endpoints can be accessed by using their private IP addresses. By using Amazon Route 53 forwarding rules, DNS queries can be resolved to the interface VPC endpoints in the shared services VPC.

#### NEW QUESTION 8

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 9

An insurance company is planning the migration of workloads from its on-premises data center to the AWS Cloud. The company requires end-to-end domain name resolution. Bi-directional DNS resolution between AWS and the existing on-premises environments must be established. The workloads will be migrated into multiple VPCs. The workloads also have dependencies on each other, and not all the workloads will be migrated at the same time.

Which solution meets these requirements?

- A. Configure a private hosted zone for each application VPC, and create the requisite record
- B. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC
- C. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver
- D. Associate the application VPC private hosted zones with the egress VPC, and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager
- E. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints.
- F. Configure a public hosted zone for each application VPC, and create the requisite record
- G. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC
- H. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver
- I. Associate the application VPC private hosted zones with the egress VPC
- J. and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager
- K. Configure the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints.

- L. Configure a private hosted zone for each application VPC, and create the requisite record
- M. Create a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC. Define Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver
- N. Associate the application VPC private hosted zones with the egress VPC and share the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager

**Answer:** A

**Explanation:**

Creating a private hosted zone for each application VPC and creating the requisite records would enable end-to-end domain name resolution for the resources. Creating a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC would enable bi-directional DNS resolution between AWS and the existing on-premises environments. Defining Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver would enable DNS queries from AWS resources to on-premises resources. Associating the application VPC private hosted zones with the egress VPC and sharing the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager would enable DNS queries among different VPCs and accounts. Configuring the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints would enable DNS queries from on-premises resources to AWS resources.

**NEW QUESTION 10**

A company is deploying third-party firewall appliances for traffic inspection and NAT capabilities in its VPC. The VPC is configured with private subnets and public subnets. The company needs to deploy the firewall appliances behind a load balancer. Which architecture will meet these requirements MOST cost-effectively?

- A. Deploy a Gateway Load Balancer with the firewall appliances as target
- B. Configure the firewall appliances with a single network interface in a private subnet
- C. Use a NAT gateway to send the traffic to the internet after inspection.
- D. Deploy a Gateway Load Balancer with the firewall appliances as target
- E. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subnet
- F. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.
- G. Deploy a Network Load Balancer with the firewall appliances as target
- H. Configure the firewall appliances with a single network interface in a private subnet
- I. Use a NAT gateway to send the traffic to the internet after inspection.
- J. Deploy a Network Load Balancer with the firewall appliances as target
- K. Configure the firewall appliances with two network interfaces: one network interface in a private subnet and another network interface in a public subnet
- L. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.

**Answer:** B

**NEW QUESTION 10**

Your company runs an application for the US market in the us-east-1 AWS region. This application uses proprietary TCP and UDP protocols on Amazon Elastic Compute Cloud (EC2) instances. End users run a real-time, front-end application on their local PCs. This front-end application knows the DNS hostname of the service. You must prepare the system for global expansion. The end users must access the application with lowest latency. How should you use AWS services to meet these requirements?

- A. Register the IP addresses of the service hosts as "A" records with latency-based routing policy in Amazon Route 53, and set a Route 53 health check for these hosts.
- B. Set the Elastic Load Balancing (ELB) load balancer in front of the hosts of the service, and register the ELB name of the main service host as an ALIAS record with a latency-based routing policy in Route 53.
- C. Set Amazon CloudFront in front of the host of the service, and register the CloudFront name of the main service as an ALIAS record in Route 53.
- D. Set the Amazon API gateway in front of the service, and register the API gateway name of the main service as an ALIAS record in Route 53.

**Answer:** B

**NEW QUESTION 15**

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 VIF
- 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 19**

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 scalin
- C. Specify the GLB in the service definitio
- 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 servic
- G. Create path-based routing rules to allow the application to target the containers that are registered in the target grou
- H. Specify the ALB in the service definitio
- 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 servic
- K. Create path-based routing rules to allow the application to target the containers that are registered in the target grou
- L. Specify the ALB in the service definitio
- 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 servic
- P. Specify the NLB in the service definitio
- Q. Create a VPC endpoint service for the NL
- 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 VP
- 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 VP
- 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
- F. Create an additional route table with only the shared services VPC attachment associated with propagated routes from the production and nonproduction VPCs.
- G. Configure a route table with the production and nonproduction VPC attachments associated with propagated routes disable
- H. 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 manages resources across VPCs in multiple AWS Regions. The company needs to connect to the resources by using its internal domain name. A network engineer needs to apply the aws.example.com DNS suffix to all resources.

What must the network engineer do to meet this requirement?

- A. Create an Amazon Route 53 private hosted zone for aws.example.com in each Region that has resource
- B. Associate the private hosted zone with that Region's VP
- C. In the appropriate private hosted zone, create DNS records for the resources in each Region.
- D. Create one Amazon Route 53 private hosted zone for aws.example.co
- E. Configure the private hosted zone to allow zone transfers with every VPC.
- F. Create one Amazon Route 53 private hosted zone for example.co
- G. Create a single resource record for aws.example.com in the private hosted zon
- H. Apply a multivalued answer routing policy to the recor
- I. Add all VPC resources as separate values in the routing policy.
- J. Create one Amazon Route 53 private hosted zone for aws.example.co
- K. Associate the private hosted zone with every VPC that has resource
- L. In the private hosted zone, create DNS records for all resources.

**Answer:** D

#### Explanation:

Creating one private hosted zone for aws.example.com and associating it with every VPC that has resources would enable DNS resolution for all resources by using their internal domain name. Creating 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 would enable private connectivity to Amazon S3 and AWS Systems Manager without using public endpoints.

#### NEW QUESTION 31

A company is deploying a non-web application on an AWS load balancer. All targets are servers located on-premises that can be accessed by using AWS Direct Connect. The company wants to ensure that the source IP addresses of clients connecting to the application are passed all the way to the end server.

How can this requirement be achieved?

- A. Use a Network Load Balancer to automatically preserve the source IP address.
- B. Use a Network Load Balancer and enable the X-Forwarded-For attribute.
- C. Use a Network Load Balancer and enable the ProxyProtocol v2 attribute.
- D. Use an Application Load Balancer to automatically preserve the source IP address in the X-Forwarded-For header.

**Answer:** C

#### Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#proxy-protocol>

#### NEW QUESTION 35

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 rules<sup>3</sup>. Creating an AWS Systems Manager Automation runbook to remediate noncompliant security groups would enable automation of the remediation process<sup>2</sup>. 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 37**

A network engineer must provide additional safeguards to protect encrypted data at Application Load Balancers (ALBs) through the use of a unique random session key.

What should the network engineer do to meet this requirement?

- A. Change the ALB security policy to a policy that supports TLS 1.2 protocol only
- B. Use AWS Key Management Service (AWS KMS) to encrypt session keys
- C. Associate an AWS WAF web ACL with the ALB
- D. and create a security rule to enforce forward secrecy (FS)
- E. Change the ALB security policy to a policy that supports forward secrecy (FS)

**Answer:** D

**NEW QUESTION 42**

A company has deployed an AWS Network Firewall firewall into a VPC. A network engineer needs to implement a solution to deliver Network Firewall flow logs to the company's Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster in the shortest possible time.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket
- B. Create an AWS Lambda function to load logs into the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster
- C. Enable Amazon Simple Notification Service (Amazon SNS) notifications on the S3 bucket to invoke the Lambda function
- D. Configure flow logs for the firewall
- E. Set the S3 bucket as the destination.
- F. Create an Amazon Kinesis Data Firehose delivery stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination
- G. Configure flow logs for the firewall Set the Kinesis Data Firehose delivery stream as the destination for the Network Firewall flow logs.
- H. Configure flow logs for the firewall
- I. Set the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination for the Network Firewall flow logs.
- J. Create an Amazon Kinesis data stream that includes the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster as the destination
- K. Configure flow logs for the firewall
- L. Set the Kinesis data stream as the destination for the Network Firewall flow logs.

**Answer:** B

**Explanation:**

<https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-analyze-aws-network-firewall-logs-usin>

**NEW QUESTION 44**

A company has two on-premises data center locations. There is a company-managed router at each data center. Each data center has a dedicated AWS Direct Connect connection to a Direct Connect gateway through a private virtual interface. The router for the first location is advertising 110 routes to the Direct Connect gateway by using BGP, and the router for the second location is advertising 60 routes to the Direct Connect gateway by using BGP. The Direct Connect gateway is attached to a company VPC through a virtual private gateway.

A network engineer receives reports that resources in the VPC are not reachable from various locations in either data center. The network engineer checks the VPC route table and sees that the routes from the first data center location are not being populated into the route table. The network engineer must resolve this issue in the most operationally efficient manner.

What should the network engineer do to meet these requirements?

- A. Remove the Direct Connect gateway, and create a new private virtual interface from each company router to the virtual private gateway of the VPC.
- B. Change the router configurations to summarize the advertised routes.
- C. Open a support ticket to increase the quota on advertised routes to the VPC route table.
- D. Create an AWS Transit Gateway
- E. Attach the transit gateway to the VPC, and connect the Direct Connect gateway to the transit gateway.

**Answer:** B

**Explanation:**

"If you advertise more than 100 routes each for IPv4 and IPv6 over the BGP session, the BGP session will go into an idle state with the BGP session

DOWN. "<https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html>

#### NEW QUESTION 49

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-group>

#### NEW QUESTION 53

A company has deployed a critical application on a fleet of Amazon EC2 instances behind an Application Load Balancer. The application must always be reachable on port 443 from the public internet. The application recently had an outage that resulted from an incorrect change to the 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 a change 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 log records to a log group in Amazon CloudWatch Log
- B. Create a CloudWatch Logs metric filter for the log group for rejected traffic
- C. Create an alarm to notify the network engineer.
- D. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture all traffic on port 443. Publish the flow log records to a log group in Amazon CloudWatch Log
- E. Create a CloudWatch Logs metric filter for the log group for all traffic
- F. Create an alarm to notify the network engineer
- G. Create a VPC Reachability Analyzer path on port 443. Specify the security group as the source
- H. Specify the EC2 instances as the destination
- I. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection
- J. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs.
- K. Create a VPC Reachability Analyzer path on port 443. Specify the internet gateway of the VPC as the source
- L. Specify the EC2 instances as the destination
- M. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection
- N. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail
- O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs.

**Answer:** C

#### NEW QUESTION 56

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 58

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 60

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 62

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 setting
- 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 65

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 AL
- 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 66

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 package
- B. Wait until the analysis has finished running to find out which EC2 instances are still listening to the old port.
- C. Enable Amazon GuardDuty
- D. Use the graphical visualizations to filter for traffic that uses the port of the old protocol
- 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 bucket
- 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 ports
- 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 VPC. 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 68

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 observed
- B. Create a new 10 Gbps dedicated connection
- 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 observed
- 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 observed
- 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 connection
- 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 A). 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 71

A company has multiple AWS accounts. Each account contains one or more VPCs. A new security guideline requires the inspection of all traffic between VPCs. The company has deployed a transit gateway that provides connectivity between all VPCs. The company also has deployed a shared services VPC with Amazon EC2 instances that include IDS services for stateful inspection. The EC2 instances are deployed across three Availability Zones. The company has set up VPC associations and routing on the transit gateway. The company has migrated a few test VPCs to the new solution for traffic inspection.

Soon after the configuration of routing, the company receives reports of intermittent connections for traffic that crosses Availability Zones.

What should a network engineer do to resolve this issue?

- A. Modify the transit gateway VPC attachment on the shared services VPC by enabling cross-Availability Zone load balancing.
- B. Modify the transit gateway VPC attachment on the shared services VPC by enabling appliance mode support.
- C. Modify the transit gateway by selecting VPN equal-cost multi-path (ECMP) routing support.
- D. Modify the transit gateway by selecting multicast support.

**Answer:** B

#### Explanation:

To resolve the issue of intermittent connections for traffic that crosses Availability Zones after configuring routing for traffic inspection between VPCs using a transit

gateway and EC2 instances with IDS services in a shared services VPC, a network engineer should modify the transit gateway VPC attachment on the shared services VPC by enabling appliance mode support (Option B). This will ensure that traffic is routed to the same EC2 instance for stateful inspection and prevent intermittent connections.

#### NEW QUESTION 76

A company is migrating an application from on premises to AWS. The company will host the application on Amazon EC2 instances that are deployed in a single VPC. During the migration period, DNS queries from the EC2 instances must be able to resolve names of on-premises servers. The migration is expected to take 3 months. After the 3-month migration period, the resolution of on-premises servers will no longer be needed. What should a network engineer do to meet these requirements with the LEAST amount of configuration?

- A. Set up an AWS Site-to-Site VPN connection between on premises and AW
- B. Deploy an Amazon Route 53 Resolver outbound endpoint in the Region that is hosting the VPC.
- C. Set up an AWS Direct Connect connection with a private VI
- D. Deploy an Amazon Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint in the Region that is hosting the VPC.
- E. Set up an AWS Client VPN connection between on premises and AW
- F. Deploy an Amazon Route 53 Resolver inbound endpoint in the VPC.
- G. Set up an AWS Direct Connect connection with a public VI
- H. Deploy an Amazon Route 53 Resolver inbound endpoint in the Region that is hosting the VP
- I. Use the IP address that is assigned to the endpoint for connectivity to the on-premises DNS servers.

**Answer:** A

#### Explanation:

Setting up an AWS Site-to-Site VPN connection between on premises and AWS would enable a secure and encrypted connection over the public internet<sup>1</sup>. Deploying an Amazon Route 53 Resolver outbound endpoint in the Region that is hosting the VPC would enable forwarding of DNS queries for on-premises servers to the on-premises DNS servers<sup>2</sup>. This would allow EC2 instances in the VPC to resolve names of on-premises servers during the migration period. After the migration period, the Route 53 Resolver outbound endpoint can be deleted with minimal configuration changes.

#### NEW QUESTION 79

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 lis
- H. Deploy this solution in all accounts.
- I. Create a new prefix lis
- J. Add all allowed IP address ranges to the prefix lis
- 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 rang
- 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 partne
- 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<sup>3</sup>. 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<sup>5</sup>. Updating security groups to use the prefix list instead of the partner IP address range would enable simplification of security group rules<sup>3</sup>. 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<sup>3</sup>.

#### NEW QUESTION 83

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 84**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### AWS-Certified-Advanced-Networking-Specialty Practice Exam Features:

- \* AWS-Certified-Advanced-Networking-Specialty Questions and Answers Updated Frequently
- \* AWS-Certified-Advanced-Networking-Specialty Practice Questions Verified by Expert Senior Certified Staff
- \* AWS-Certified-Advanced-Networking-Specialty Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* AWS-Certified-Advanced-Networking-Specialty Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AWS-Certified-Advanced-Networking-Specialty Practice Test Here](#)