Amazon

Exam Questions AWS-Certified-Advanced-Networking-Specialty

Amazon AWS Certified Advanced Networking - Specialty

NEW QUESTION 1
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 2
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 3
A company is using custom DNS servers that run BIND for name resolution in its VPCs. The VPCs are deployed across multiple AWS accounts that are part of the same organization in AWS Organizations. All the VPCs are connected to a transit gateway. The BIND servers are running in a central VPC and are configured to forward all queries for an on-premises DNS domain to DNS servers that are hosted in an on-premises data center. To ensure that all the VPCs use the custom DNS servers, a network engineer has configured a VPC DHCP options set in all the VPCs that specifies the custom DNS servers to be used as domain name servers.
Multiple development teams in the company want to use Amazon Elastic File System (Amazon EFS). A development team has created a new EFS file system but cannot mount the file system to one of its Amazon EC2 instances. The network engineer discovers that the EC2 instance cannot resolve the IP address for the EFS mount point fs-33444567d.efs.us-east-1.amazonaws.com. The network engineer needs to implement a solution so that development teams throughout the organization can mount EFS file systems.
Which combination of steps will meet these requirements? (Choose two.)

A. Configure the BIND DNS servers in the central VPC to forward queries forefs.us-east-1.amazonaws.com to the Amazon provided DNS server (169.254.169.253).
B. Create an Amazon Route 53 Resolver outbound endpoint in the central VP
C. Update all the VPC DHCP options sets to use AmazonProvidedDNS for name resolution.
D. Create an Amazon Route 53 Resolver inbound endpoint in the central VPUpdate all the VPC DHCP options sets to use the Route 53 Resolver inbound endpoint in the central VPC for name resolution.
E. Create an Amazon Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS server
F. Share the rule with the organization by using AWS Resource Access Manager (AWS RAM). Associate the rule with all the VPCs.
G. Create an Amazon Route 53 private hosted zone for the efs.us-east-1.amazonaws.com domain.Associate the private hosted zone with the VPC where the EC2 instance is deploye
H. Create an A record for fs-33444567d.efs.us-east-1.amazonaws.com in the private hosted zon
I. Configure the A record to return the mount target of the EFS mount point.

Answer: BD

Explanation:

Option B suggests using Amazon Route 53 Resolver outbound endpoint, which would replace the existing BIND DNS servers with the AmazonProvidedDNS for name resolution. However, the scenario specifically mentions that the company is using custom DNS servers that run BIND for name resolution in its VPCs, so this solution would not work. Option D suggests creating a Route 53 Resolver rule to forward queries for the on-premises domain to the on-premises DNS servers, which would not address the issue of resolving the EFS mount point. The problem is not with resolving queries for the on-premises domain, but rather with resolving the IP address for the EFS mount point.

NEW QUESTION 4
An IoT company sells hardware sensor modules that periodically send out temperature, humidity, pressure, and location data through the MQTT messaging protocol. The hardware sensor modules send this data to the company's on-premises MQTT brokers that run on Linux servers behind a load balancer. The hardware sensor modules have been hardcoded with public IP addresses to reach the brokers.
The company is growing and is acquiring customers across the world. The existing solution can no longer scale and is introducing additional latency because of the company's global presence. As a result, the company decides to migrate its entire infrastructure from on premises to the AWS Cloud. The company needs to migrate without reconfiguring the hardware sensor modules that are already deployed across the world. The solution also must minimize latency.
The company migrates the MQTT brokers to run on Amazon EC2 instances. What should the company do next to meet these requirements?

A. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listener
B. Use Bring Your Own IP (BYOIP) from the on-premises network with the NLB.
C. Place the EC2 instances behind a Network Load Balancer (NLB). Configure TCP listener
D. Create an AWS Global Accelerator accelerator in front of the NLUse Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator.
E. Place the EC2 instances behind an Application Load Balancer (ALB). Configure TCP listener
F. Create an AWS Global Accelerator accelerator in front of the AL
G. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator
H. Place the EC2 instances behind an Amazon CloudFront distributio
I. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

Answer: B

NEW QUESTION 5
A network engineer needs to standardize a company's approach to centralizing and managing interface VPC endpoints for private communication with AWS services. The company uses AWS Transit Gateway for inter-VPC connectivity between AWS accounts through a hub-and-spoke model. The company's network services team must manage all Amazon Route 53 zones and interface endpoints within a shared services AWS account. The company wants to use thiscentralized model to provide AWS resources with access to AWS Key Management Service (AWS KMS) without sending traffic over the public internet.
What should the network engineer do to meet these requirements?

A. In the shared services account, create an interface endpoint for AWS KM
B. Modify the interface endpoint by disabling the private DNS nam
C. Create a private hosted zone in the shared services account with an alias record that points to the interface endpoin
D. Associate the private hosted zone with the spoke VPCs in each AWS account.
E. In the shared services account, create an interface endpoint for AWS KM
F. Modify the interface endpoint by disabling the private DNS nam
G. Create a private hosted zone in each spoke AWS account with an alias record that points to the interface endpoin
H. Associate each private hosted zone with the shared services AWS account.
I. In each spoke AWS account, create an interface endpoint for AWS KM
J. Modify each interface endpoint by disabling the private DNS nam
K. Create a private hosted zone in each spoke AWS account with an alias record that points to each interface endpoin
L. Associate each private hosted zone with the shared services AWS account.
M. In each spoke AWS account, create an interface endpoint for AWS KM
N. Modify each interface endpoint by disabling the private DNS nam
O. Create a private hosted zone in the shared services account with an alias record that points to each interface endpoin
P. Associate the private hosted zone with the spoke VPCs in each AWS account.

Answer: A

NEW QUESTION 6
A company has deployed a web application on AWS. The web application uses an Application Load Balancer (ALB) across multiple Availability Zones. The targets of the ALB are AWS Lambda functions. The web application also uses Amazon CloudWatch metrics for monitoring.
Users report that parts of the web application are not loading properly. A network engineer needs to troubleshoot the problem. The network engineer enables access logging for the ALB.
What should the network engineer do next to determine which errors the ALB is receiving?

A. Send the logs to Amazon CloudWatch Log
B. Review the ALB logs in CloudWatch Insights to determine which error messages the ALB is receiving.
C. Configure the Amazon S3 bucket destinatio
D. Use Amazon Athena to determine which error messages the ALB is receiving.
E. Configure the Amazon S3 bucket destinatio
F. After Amazon CloudWatch Logs pulls the ALB logs from the S3 bucket automatically, review the logs in CloudWatch Logs to determine which error messages the ALB is receiving.
G. Send the logs to Amazon CloudWatch Log
H. Use the Amazon Athena CloudWatch Connector todetermine which error messages the ALB is receiving.

Answer: B

Explanation:
Access logs is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logs for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify as compressed files. You can disable access logs at any time.https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html

NEW QUESTION 7
All IP addresses within a 10.0.0.0/16 VPC are fully utilized with application servers across two Availability Zones. The application servers need to send frequent

UDP probes to a single central authentication server on the Internet to confirm that is running up-to-date packages. The network is designed for application servers to use a single NAT gateway for internal access. Testing reveals that a few of the servers are unable to communicate with the authentication server.

A. The NAT gateway does not support UDP traffic.
B. The authentication server is not accepting traffic.
C. The NAT gateway cannot allocate more ports.
D. The NAT gateway is launched in a private subnet.

Answer: C

Explanation:
Ref:https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html
"A NAT gateway can support up to 55,000 simultaneous connections to each unique destination. This limit also applies if you create approximately 900 connections per second to a single destination (about 55,000 connections per minute). If the destination IP address, the destination port, or the protocol (TCP/UDP/ICMP) changes, you can create an additional 55,000 connections. For more than 55,000 connections, there is an increased chance of connection errors due to port allocation errors. These errors can be monitored by viewing the ErrorPortAllocation CloudWatch metric for your NAT gateway. For more information, see Monitoring NAT Gateways Using Amazon CloudWatch."

NEW QUESTION 8
Your organization has a newly installed 1-Gbps AWS Direct Connect connection. You order the cross-connect from the Direct Connect location provider to the port on your router in the same facility. To enable the use of your first virtual interface, your router must be configured appropriately.
What are the minimum requirements for your router?

A. 1-Gbps Multi Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
B. 1-Gbps Single Mode Fiber Interface, 802.1Q VLAN, Peer IP Address, BGP Session with MD5.
C. IPsec Parameters, Pre-Shared key, Peer IP Address, BGP Session with MD5
D. BGP Session with MD5, 802.1Q VLAN, Route-Map, Prefix List, IPsec encrypted GRE Tunnel

Answer: B

NEW QUESTION 9
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 subne
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 subne
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 subne
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 subne
L. Use the NAT functionality on the firewall appliances to send the traffic to the internet after inspection.

Answer: B

NEW QUESTION 10
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 templat
B. Define the primary network interface to be created in one of the private subnet
C. For the second network interface, select one of the public subnet
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 templat
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 launchin
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 interfac
J. Use the user data option to run a cloud-init script to allocate a second network interface and to associate anElastic 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 company is planning a migration of its critical workloads from an on-premises data center to Amazon EC2 instances. The plan includes a new 10 Gbps AWS Direct Connect dedicated connection from the on-premises data center to a VPC that is attached to a transit gateway. The migration must occur over encrypted paths between the on-premises data center and the AWS Cloud.
Which solution will meet these requirements while providing the HIGHEST throughput?

A. Configure a public VIF on the Direct Connect connectio
B. Configure an AWS Site-to-Site VPN connection to the transit gateway as a VPN attachment.
C. Configure a transit VIF on the Direct Connect connectio
D. Configure an IPsec VPN connection to an EC2 instance that is running third-party VPN software.
E. Configure MACsec for the Direct Connect connectio
F. Configure a transit VIF to a Direct Connect gateway that is associated with the transit gateway.
G. Configure a public VIF on the Direct Connect connectio
H. Configure two AWS Site-to-Site VPN connections to the transit gatewa
I. Enable equal-cost multi-path (ECMP) routing.

Answer: C

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/adding-macsec-security-to-aws-direct-connect-c

NEW QUESTION 11
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 14
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 18
A company has deployed a software-defined WAN (SD-WAN) solution to interconnect all of its offices. The company is migrating workloads to AWS 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. According to 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 applianc
B. Add routes that are more specific to point to the primary SD-WAN virtual appliance.
C. Configure the BGP community tag 7224:7300 on the primary SD-WAN virtual appliance for BGP routes toward the transit gateway.
D. Configure the AS_PATH prepend attribute on the secondary SD-WAN virtual appliance for BGP routes toward the transit gateway.
E. Disable equal-cost multi-path (ECMP) routing on the transit gateway for Transit Gateway Connect.

Answer: A

NEW QUESTION 21
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 VPCreate 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 disable
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 26
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 multivalue 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 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 configuratio
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 configuratio
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 configuratio
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 configuratio
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 29
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 bucke
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 bucke
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 30
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 bucke
B. Create an AWS Lambda function to load logs into the Amazon OpenSearch Service (Amazon Elasticsearch Service) cluste
C. Enable Amazon Simple Notification Service (Amazon SNS) notifications on the S3 bucket to invoke the Lambda functio
D. Configure flow logs for the firewal
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 destinatio
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 firewal
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 destinatio
K. Configure flow logs for the firewal
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 33
A global company operates all its non-production environments out of three AWS Regions: eu-west-1,
us-east-1, and us-west-1. The company hosts all its production workloads in two on-premises data centers. The company has 60 AWS accounts and each account has two VPCs in each Region. Each VPC has a virtual private gateway where two VPN connections terminate for resilient connectivity to the data centers. The company has 360 VPN tunnels to each data center, resulting in high management overhead. The total VPN throughput for each Region is 500 Mbps.
The company wants to migrate the production environments to AWS. The company needs a solution that will simplify the network architecture and allow for future growth. The production environments will generate an additional 2 Gbps of traffic per Region back to the data centers. This traffic will increase over time.
Which solution will meet these requirements?

A. Set up an AWS Direct Connect connection from each data center to AWS in each Regio
B. Create and attach private VIFs to a single Direct Connect gatewa
C. Attach the Direct Connect gateway to all the VPC
D. Remove the existing VPN connections that are attached directly to the virtual private gateways.
E. Create a single transit gateway with VPN connections from each data cente
F. Share the transit gateway with each account by using AWS Resource Access Manager (AWS RAM). Attach the transit gateway to each VP
G. Remove the existing VPN connections that are attached directly to the virtual private gateways.
H. Create a transit gateway in each Region with multiple newly commissioned VPN connections from each data cente
I. Share the transit gateways with each account by using AWS Resource Access Manager (AWS RAM). In each Region, attach the transit gateway to each VPRemove the existing VPN connections that are attached directly to the virtual private gateways.
J. Peer all the VPCs in each Region to a new VPC in each Region that will function as a centralized transit VP
K. Create new VPN connections from each data center to the transit VPC
L. Terminate the original VPN connections that are attached to all the original VPC
M. Retain the new VPN connection to the new transit VPC in each Region.

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 41
An organization launched an IPv6-only web portal to support IPv6-native mobile clients. Front-end instances launch in an Amazon VPC associated with an appropriate IPv6 CIDR. The VPC IPv4 CIDR is fully utilized. A single subnet exists in each of two Availability Zones with appropriately configured IPv6 CIDR associations. Auto Scaling is properly configured, and no Elastic Load Balancing is used.
Customers say the service is unavailable during peak load times. The network engineer attempts to launch an instance manually and receives the following message: "There are not enough free addresses in subnet 'subnet-12345677' to satisfy the requested number of instances."

What action will resolve the availability problem?

A. Create a new subnet using a VPC secondary IPv6 CIDR, and associate an IPv6 CID
B. Include the new subnet in the Auto Scaling group.
C. Create a new subnet using a VPC secondary IPv4 CIDR, and associate an IPv6 CID
D. Include the new subnet in the Auto Scaling group.
E. Resize the IPv6 CIDR on each of the existing subnet
F. Modify the Auto Scaling group maximum number of instances.
G. Add a secondary IPv4 CIDR to the Amazon VP
H. Assign secondary IPv4 address space to each of theexisting subnets.

Answer: B

NEW QUESTION 45
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 49
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 payloa
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 payloa
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 payloa
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 payloa
H. Configure the tool to insert a customized header to inform the web application of an authenticated customer request.

Answer: C

NEW QUESTION 53
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 58
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 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 63
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 68
A media company is implementing a news website for a global audience. The website uses Amazon CloudFront as its content delivery network. The backend runs on Amazon EC2 Windows instances behind an Application Load Balancer (ALB). The instances are part of an Auto Scaling group. The company's customers access the website by using service example com as the CloudFront custom domain name. The CloudFront origin points to an ALB that uses service-alb.example.com as the domain name.
The company's security policy requires the traffic to be encrypted in transit at all times between the users and the backend.
Which combination of changes must the company make to meet this security requirement? (Choose three.)

A. Create a self-signed certificate for service.example.co
B. Import the certificate into AWS Certificate Manager (ACM). Configure CloudFront to use this imported SSL/TLS certificat
C. Change the default behavior to redirect HTTP to HTTPS.
D. Create a certificate for service.example.com by using AWS Certificate Manager (ACM). Configure CloudFront to use this custom SSL/TLS certificat
E. Change the default behavior to redirect HTTP to HTTPS.
F. Create a certificate with any domain name by using AWS Certificate Manager (ACM) for the EC2 instance
G. Configure the backend to use this certificate for its HTTPS listene
H. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its target
I. Attach the existing Auto Scaling group to this new target group.
J. Create a public certificate from a third-party certificate provider with any domain name for the EC2 instance
K. Configure the backend to use this certificate for its HTTPS listene
L. Specify the instance target type during the creation of a new target group that uses the HTTPS protocol for its target
M. Attach the existing Auto Scaling group to this new target group.
N. Create a certificate for service-alb.example.com by using AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the service-alb.example.com ACM certificat
O. Modify the CloudFront origin to use the HTTPS protocol onl
P. Delete the HTTPlistener on the ALB.
Q. Create a self-signed certificate for service-alb.example.co
R. Import the certificate into AWS Certificate Manager (ACM). On the ALB add a new HTTPS listener that uses the new target group and the imported service-alb.example.com ACM certificat
S. Modify the CloudFront origin to use the HTTPS protocol onl
T. Delete the HTTP listener on the ALB.

Answer: BDE

NEW QUESTION 73
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 internet1. 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 servers2. 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 78
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 81
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 quot
D. Use AWS Site-to-Site VPNto 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 gatewa
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 transitgateway, 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 82
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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