



Amazon

Exam Questions AWS-Certified-Advanced-Networking-Specialty

Amazon AWS Certified Advanced Networking - Specialty

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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

- A. Inbound; Protocol tcp; Source [Instance's EIP]; Destination 169.254.169.254
- B. Inbound; Protocol tcp; Destination 169.254.169.254; Destination port 80
- C. Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 80
- D. Outbound; Protocol tcp; Destination 169.254.169.254; Destination port 443

Answer: C

Explanation:

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

To view all categories of instance metadata from within a running instance, use the following URI.

<http://169.254.169.254/latest/meta-data/>

NEW QUESTION 2

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 for `efs.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 VPC
- 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 VPC
- 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 deployed
- H. Create an A record for `fs-33444567d.efs.us-east-1.amazonaws.com` in the private hosted zone
- 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 3

A company has an AWS Site-to-Site VPN connection between its existing VPC and on-premises network. The default DHCP options set is associated with the VPC. The company has an application that is running on an Amazon Linux 2 Amazon EC2 instance in the VPC. The application must retrieve an Amazon RDS database secret that is stored in AWS Secrets Manager through a private VPC endpoint. An on-premises application provides internal RESTful API service that can be reached by URL (<https://api.example.internal>). Two on-premises Windows DNS servers provide internal DNS resolution.

The application on the EC2 instance needs to call the internal API service that is deployed in the on-premises environment. When the application on the EC2 instance attempts to call the internal API service by referring to the hostname that is assigned to the service, the call fails. When a network engineer tests the API service call from the same EC2 instance by using the API service's IP address, the call is successful.

What should the network engineer do to resolve this issue and prevent the same problem from affecting other resources in the VPC?

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

Answer: B

Explanation:

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

NEW QUESTION 4

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

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 5

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

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

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

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

Answer: C

NEW QUESTION 6

A company is planning to use Amazon S3 to archive financial data. The data is currently stored in an on-premises data center. The company uses AWS Direct Connect with a Direct Connect gateway and a transit gateway to connect to the on-premises data center. The data cannot be transported over the public internet and must be encrypted in transit.

Which solution will meet these requirements?

- A. Create a Direct Connect public VI
- B. Set up an IPsec VPN connection over the public VIF to access Amazon S3. Use HTTPS for communication.
- C. Create an IPsec VPN connection over the transit VI
- D. Create a VPC and attach the VPC to the transit gatewa
- E. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.
- F. Create a VPC and attach the VPC to the transit gatewa
- G. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.
- H. Create a Direct Connect public VI
- I. Set up an IPsec VPN connection over the public VIF to the transit gatewa
- J. Create an attachment for Amazon S3. Use HTTPS for communication.

Answer: B

Explanation:

<https://docs.aws.amazon.com/vpn/latest/s2svpn/private-ip-dx.html>

An IPsec VPN connection over the transit VIF can encrypt traffic between the on-premises network and AWS without using public IP addresses or the internet2. A VPC endpoint for Amazon S3 can enable private access to S3 buckets within the same region. HTTPS can provide additional encryption for communication.

NEW QUESTION 7

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

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

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

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

Answer: C

NEW QUESTION 8

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. 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 gatewa
- E. Advertise the VPC-B CIDR block under the allowed prefixes.
- F. Configure another transit VIF on the Direct Connect connection and associate TGW-
- G. Advertise the VPC-B CIDR block under the allowed prefixes.
- H. Configure inter-Region transit gateway peering between TGW-A and TGW-
- 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 9

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 10

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 10

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 14

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 17

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 20

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

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

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

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

Answer: AB

Explanation:

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

NEW QUESTION 21

A company has two AWS accounts one for Production and one for Connectivity. A network engineer needs to connect the Production account VPC to a transit gateway in the Connectivity account. The feature to auto accept shared attachments is not enabled on the transit gateway.

Which set of steps should the network engineer follow in each AWS account to meet these requirements?

- A. * 1. In the Production account: Create a resource share in AWS Resource Access Manager for the transit gateway
- B. Provide the Connectivity account ID
- C. Enable the feature to allow external accounts* 2. In the Connectivity account: Accept the resource.* 3. In the Connectivity account: Create an attachment to the VPC subnets.* 4. In the Production account: Accept the attachment
- D. Associate a route table with the attachment.
- E. * 1. In the Production account: Create a resource share in AWS Resource Access Manager for the VPC subnet
- F. Provide the Connectivity account ID
- G. Enable the feature to allow external accounts.* 2. In the Connectivity account: Accept the resource.* 3. In the Production account: Create an attachment on the transit gateway to the VPC subnets.* 4. In the Connectivity account: Accept the attachment
- H. Associate a route table with the attachment.
- I. * 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the VPC subnet
- J. Provide the Production account ID
- K. Enable the feature to allow external accounts.* 2. In the Production account: Accept the resource.* 3. In the Connectivity account: Create an attachment on the transit gateway to the VPC subnets.* 4. In the Production account: Accept the attachment
- L. Associate a route table with the attachment.
- M. * 1. In the Connectivity account: Create a resource share in AWS Resource Access Manager for the transit gateway
- N. Provide the Production account ID Enable the feature to allow external accounts.* 2. In the Production account: Accept the resource.* 3. In the Production account: Create an attachment to the VPC subnets.* 4. In the Connectivity account: Accept the attachment
- O. Associate a route table with the attachment.

Answer: A

Explanation:

step 1: In the Production account, create a resource share in AWS Resource Access Manager for the transit gateway and provide the Connectivity account ID. Enabling the feature to allow external accounts is also required to share resources between accounts. Step 2: In the Connectivity account, accept the shared resource. This action will allow the Production account to use the transit gateway in the Connectivity account. Step 3: In the Connectivity account, create an attachment to the VPC subnets. This attachment will enable communication between the VPC in the Production account and the transit gateway in the Connectivity account. Step 4: In the Production account, accept the attachment and associate a route table with the attachment. This will enable the VPC to route traffic through the transit gateway to other resources in the Connectivity account.

NEW QUESTION 23

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 connection
- 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 connection
- D. Configure an IPsec VPN connection to an EC2 instance that is running third-party VPN software.
- E. Configure MACsec for the Direct Connect connection
- 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 connection
- H. Configure two AWS Site-to-Site VPN connections to the transit gateway
- 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 26

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 31

A company has expanded its network to the AWS Cloud by using a hybrid architecture with multiple AWS accounts. The company has set up a shared AWS account for the connection to its on-premises data centers and the company offices. The workloads consist of private web-based services for internal use. These services run in different AWS accounts. Office-based employees consume these services by using a DNS name in an on-premises DNS zone that is named example.internal.

The process to register a new service that runs on AWS requires a manual and complicated change request to the internal DNS. The process involves many teams.

The company wants to update the DNS registration process by giving the service creators access that will allow them to register their DNS records. A network engineer must design a solution that will achieve this goal. The solution must maximize cost-effectiveness and must require the least possible number of configuration changes.

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

- A. Create a record for each service in its local private hosted zone (serviceA.account1.aws.example.internal). Provide this DNS record to the employees who need access.
- B. Create an Amazon Route 53 Resolver inbound endpoint in the shared account VPC
- C. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS server
- D. Set the forwarding IP addresses to the inbound endpoint's IP addresses that were created.
- E. Create an Amazon Route 53 Resolver rule to forward any queries made to onprem.example.internal to the on-premises DNS servers.
- F. Create an Amazon Route 53 private hosted zone named aws.example.internal in the shared AWS account to resolve queries for this domain.
- G. Launch two Amazon EC2 instances in the shared AWS account
- H. Install BIND on each instance
- I. Create a DNS conditional forwarder on each BIND server to forward queries for each subdomain under aws.example.internal to the appropriate private hosted zone in each AWS account
- J. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS server
- K. Set the forwarding IP addresses to the IP addresses of the BIND servers.
- L. Create a private hosted zone in the shared AWS account for each account that runs the service. Configure the private hosted zone to contain aws.example.internal in the domain (account1.aws.example.internal). Associate the private hosted zone with the VPC that runs the service and the shared account VPC.

Answer: ABD

Explanation:

To meet the requirements of updating the DNS registration process while maximizing cost-effectiveness and minimizing configuration changes, the network engineer should take the following steps:

- Create an Amazon Route 53 Resolver inbound endpoint in the shared account VPC. Create a conditional forwarder for a domain named aws.example.internal on the on-premises DNS servers. Set the forwarding IP addresses to the inbound endpoint's IP addresses that were created (Option B).
- Create an Amazon Route 53 private hosted zone named aws.example.internal in the shared AWS account to resolve queries for this domain (Option D).
- Create a record for each service in its local private hosted zone (serviceA.account1.aws.example.internal). Provide this DNS record to the employees who need access (Option A).

These steps will allow service creators to register their DNS records while keeping costs low and minimizing configuration changes.

NEW QUESTION 36

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

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

- A. Update the DNS Firewall VPC configuration to disable fail open for the VPC.
- B. Update the DNS Firewall VPC configuration to enable fail open for the VPC.
- C. Create a new DHCP options set with parameter dns_firewall_fail_open=false

- D. Associate the new DHCP options set with the VPC.
- E. Create a new DHCP options set with parameter dns_firewall_fail_open=tru
- F. Associate the new DHCP options set with the VPC.

Answer: B

NEW QUESTION 41

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

Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 44

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 46

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 51

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-using-Amazon-Kinesis-Data-Firehose/>

NEW QUESTION 52

A company has deployed its AWS environment in a single AWS Region. The environment consists of a few hundred application VPCs, a shared services VPC, and a VPN connection to the company's on-premises environment. A network engineer needs to implement a transit gateway with the following requirements:

- Application VPCs must be isolated from each other.
 - Bidirectional communication must be allowed between the application VPCs and the on-premises network.
 - Bidirectional communication must be allowed between the application VPCs and the shared services VPC. The network engineer creates the transit gateway with options disabled for default route table association and default route table propagation. The network engineer also creates the VPN attachment for the on-premises network and creates the VPC attachments for the application VPCs and the shared services VPC.
- The network engineer must meet all the requirements for the transit gateway by designing a solution that needs the least number of transit gateway route tables. Which combination of actions should the network engineer perform to accomplish this goal?(Choose two.)

- A. Configure a separate transit gateway route table for on-premise.
- B. Associate the VPN attachment with this transit gateway route table.
- C. Propagate all application VPC attachments to this transit gateway route table.
- D. Configure a separate transit gateway route table for each application VPC.
- E. Associate each application VPC attachment with its respective transit gateway route table.
- F. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.
- G. Configure a separate transit gateway route table for all application VPCs.
- H. Associate all application VPCs with this transit gateway route table.
- I. Propagate the shared services VPC attachment and the VPN attachment to this transit gateway route table.
- J. Configure a separate transit gateway route table for the shared services VPC.
- K. Associate the shared services VPC attachment with this transit gateway route table.
- L. Propagate all application VPC attachments to this transit gateway route table.
- M. Configure a separate transit gateway route table for on-premises and the shared services VPC.
- N. Associate the VPN attachment and the shared services VPC attachment with this transit gateway route table.
- O. Propagate all application VPC attachments to this transit gateway route table.

Answer: BD

NEW QUESTION 55

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-requirements>

NEW QUESTION 60

A software company offers a software-as-a-service (SaaS) accounting application that is hosted in the AWS Cloud. The application requires connectivity to the company's on-premises network. The company has two redundant 10 GB AWS Direct Connect connections between AWS and its on-premises network to accommodate the growing demand for the application. The company already has encryption between its on-premises network and the colocation. The company needs to encrypt traffic between AWS and the edge routers in the colocation within the next few months. The company must maintain its current bandwidth. What should a network engineer do to meet these requirements with the LEAST operational overhead?

- A. Deploy a new public VIF with encryption on the existing Direct Connect connection.
- B. Reroute traffic through the new public VIF.

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

Answer: C

NEW QUESTION 61

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 65

A company is planning to deploy many software-defined WAN (SD-WAN) sites. The company is using AWS Transit Gateway and has deployed a transit gateway in the required AWS Region. A network engineer needs to deploy the SD-WAN hub virtual appliance into a VPC that is connected to the transit gateway. The solution must support at least 5 Gbps of throughput from the SD-WAN hub virtual appliance to other VPCs that are attached to the transit gateway. Which solution will meet these requirements?

- A. Create a new VPC for the SD-WAN hub virtual appliance
- B. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway
- C. Configure BGP over the IPsec VPN connections
- D. Assign a new CIDR block to the transit gateway
- E. Create a new VPC for the SD-WAN hub virtual appliance
- F. Attach the new VPC to the transit gateway with a VPC attachment
- G. Add a transit gateway Connect attachment
- H. Create a Connect peer and specify the GRE and BGP parameter
- I. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.
- J. Create a new VPC for the SD-WAN hub virtual appliance
- K. Attach the new VPC to the transit gateway with a VPC attachment
- L. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gateway
- M. Configure BGP over the IPsec VPN connections.
- N. Assign a new CIDR block to the transit gateway
- O. Create a new VPC for the SD-WAN hub virtual appliance
- P. Attach the new VPC to the transit gateway with a VPC attachment
- Q. Add a transit gateway Connect attachment
- R. Create a Connect peer and specify the VXLAN and BGP parameter
- S. Create a route in the appropriate VPC for the SD-WAN hub virtual appliance to route to the transit gateway.

Answer: D

NEW QUESTION 66

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 67

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 network
- I. Configure data center routers to make routing decisions based on the BGP communities received.

Answer: AD

NEW QUESTION 69

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 73

A Network Engineer is provisioning a subnet for a load balancer that will sit in front of a fleet of application servers in a private subnet. There is limited IP space left in the VPC CIDR. The application has few users now but is expected to grow quickly to millions of users.

What design will use the LEAST amount of IP space, while allowing for this growth?

- A. Use two /29 subnets for an Application Load Balancer in different Availability Zones.
- B. Use one /29 subnet for the Network Load Balance
- C. Add another VPC CIDR to the VPC to allow for future growth.
- D. Use two /28 subnets for a Network Load Balancer in different Availability Zones.
- E. Use one /28 subnet for an Application Load Balance
- F. Add another VPC CIDR to the VPC to allow for future growth.

Answer: C

NEW QUESTION 76

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 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 77

A bank built a new version of its banking application in AWS using containers that connect to an on-premises database over VPN connection. This application version requires users to also update their client application. The bank plans to deprecate the earlier client version. However, the company wants to keep supporting earlier clients through their on-premises version of the application to serve a small portion of the customers who haven't yet upgraded. What design will allow the company to serve both newer and earlier clients in the MOST efficient way?

- A. Use an Amazon Route 53 multivalue answer routing policy to route older client traffic to the on-premises application version and the rest of the traffic to the new AWS based version.
- B. Use a Classic Load Balancer for the new application
- C. Route all traffic to the new application by using an Elastic Load Balancing (ELB) load balancer
- D. Define a user-agent-based rule on the backend servers to redirect earlier clients to the on-premises application.
- E. Use an Application Load Balancer for the new application
- F. Register both the new and earlier applications as separate target groups and use path-based routing to route traffic based on the application version.
- G. Use an Application Load Balancer for the new application
- H. Register both the new and earlier application backends as separate target groups
- I. Use header-based routing to route traffic based on the application version.

Answer: D

NEW QUESTION 80

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

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

Answer: D

Explanation:

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

NEW QUESTION 85

A company is deploying a new application on AWS. The application uses dynamic multicasting. The company has five VPCs that are all attached to a transit gateway. Amazon EC2 instances in each VPC need to be able to register dynamically to receive a multicast transmission. How should a network engineer configure the AWS resources to meet these requirements?

- A. Create a static source multicast domain within the transit gateway
- B. Associate the VPCs and applicable subnets with the multicast domain
- C. Register the multicast senders' network interface with the multicast domain

- D. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
- E. Create a static source multicast domain within the transit gateway
- F. Associate the VPCs and applicable subnets with the multicast domain
- G. Register the multicast senders' network interface with the multicast domain
- H. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.
- I. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain
- J. Register the multicast senders' network interface with the multicast domain
- K. Adjust the network ACLs to allow UDP traffic from the source to all receivers and to allow UDP traffic that is sent to the multicast group address.
- L. Create an Internet Group Management Protocol (IGMP) multicast domain within the transit gateway. Associate the VPCs and applicable subnets with the multicast domain
- M. Register the multicast senders' network interface with the multicast domain
- N. Adjust the network ACLs to allow TCP traffic from the source to all receivers and to allow TCP traffic that is sent to the multicast group address.

Answer: C

NEW QUESTION 89

A company uses a hybrid architecture and has an AWS Direct Connect connection between its on-premises data center and AWS. The company has production applications that run in the on-premises data center. The company also has production applications that run in a VPC. The applications that run in the on-premises data center need to communicate with the applications that run in the VPC. The company is using corp.example.com as the domain name for the on-premises resources and is using an Amazon Route 53 private hosted zone for aws.example.com to host the VPC resources.

The company is using an open-source recursive DNS resolver in a VPC subnet and is using a DNS resolver in the on-premises data center. The company's on-premises DNS resolver has a forwarder that directs requests for the aws.example.com domain name to the DNS resolver in the VPC. The DNS resolver in the VPC has a forwarder that directs requests for the corp.example.com domain name to the DNS resolver in the on-premises data center. The company has decided to replace the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints.

Which combination of steps should a network engineer take to make this replacement? (Choose three.)

- A. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the outbound endpoint.
- B. Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.
- C. Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint.
- D. Create a Route 53 Resolver rule to forward aws.example.com domain queries to the IP addresses of the inbound endpoint.
- E. Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver.
- F. Configure the on-premises DNS resolver to forward aws.example.com queries to the IP addresses of the outbound endpoint.

Answer: BCE

Explanation:

To replace the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints in a hybrid architecture where on-premises applications need to communicate with applications running in a VPC, a network engineer should take the following steps:

- Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint. (Option C)
- Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint. (Option B)
- Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver. (Option E)

These steps will allow for seamless replacement of the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints and enable communication between on-premises and VPC applications.

NEW QUESTION 93

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

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

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

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

Answer: D

NEW QUESTION 94

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 96

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

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

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 97

.....

Relate Links

100% Pass Your AWS-Certified-Advanced-Networking-Specialty Exam with ExamBible Prep Materials

<https://www.exambible.com/AWS-Certified-Advanced-Networking-Specialty-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>