



Amazon-Web-Services

Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam

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

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 NLB. Use 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 ALB
- G. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator
- H. Place the EC2 instances behind an Amazon CloudFront distribution
- I. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

Answer: B

NEW QUESTION 2

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

Answer: A

NEW QUESTION 3

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 VIF
- 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 VIF
- D. Create a VPC and attach the VPC to the transit gateway
- 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 gateway
- G. In the VPC, provision an interface VPC endpoint for Amazon S3. Use HTTPS for communication.
- H. Create a Direct Connect public VIF
- I. Set up an IPsec VPN connection over the public VIF to the transit gateway
- 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 internet. 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 4

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 gateway
- E. Advertise the VPC-B CIDR block under the allowed prefixes.
- F. Configure another transit VIF on the Direct Connect connection and associate TGW-
- 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 5

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 6

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

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

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

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

Answer: C

NEW QUESTION 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

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

Which solution meets these requirements?

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

Answer: A

Explanation:

Creating a private hosted zone for each application VPC and creating the requisite records would enable end-to-end domain name resolution for the resources. Creating a set of Amazon Route 53 Resolver inbound and outbound endpoints in an egress VPC would enable bi-directional DNS resolution between AWS and the existing on-premises environments. Defining Route 53 Resolver rules to forward requests for the on-premises domains to the on-premises DNS resolver would enable DNS queries from AWS resources to on-premises resources. Associating the application VPC private hosted zones with the egress VPC and sharing the Route 53 Resolver rules with the application accounts by using AWS Resource Access Manager would enable DNS queries among different VPCs and accounts. Configuring the on-premises DNS servers to forward the cloud domains to the Route 53 inbound endpoints would enable DNS queries from on-premises resources to AWS resources1.

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 template
- 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 template
- F. Use the user data option to run a cloud-init script after boot to attach the second network interface from a subnet with auto-assign public IP addressing enabled.
- G. Create an AWS Lambda function to run as a lifecycle hook of the Auto Scaling group when an instance is launched
- H. In the Lambda function, assign a network interface to an AWS Global Accelerator endpoint.
- I. During creation of the Auto Scaling group, select subnets for the primary network interface
- J. Use the user data option to run a cloud-init script to allocate a second network interface and to associate an Elastic IP address from the BYOIP pool.

Answer: D

Explanation:

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

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

NEW QUESTION 10

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

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

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

- A. Deploy the SaaS service endpoint behind a Network Load Balancer.
- B. Configure an endpoint service, and grant the customers permission to create a connection to the endpoint service.
- C. Deploy the SaaS service endpoint behind an Application Load Balancer.
- D. Configure a VPC peering connection to the customer VPC
- E. Route traffic through NAT gateways.
- F. Deploy an AWS Transit Gateway, and connect the SaaS VPC to 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 12

Your company runs an application for the US market in the us-east-1 AWS region. This application uses proprietary TCP and UDP protocols on Amazon Elastic Compute Cloud (EC2) instances. End users run a

real-time, front-end application on their local PCs. This front-end application knows the DNS hostname of the service.

You must prepare the system for global expansion. The end users must access the application with lowest latency.

How should you use AWS services to meet these requirements?

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

Answer: B

NEW QUESTION 17

A network engineer needs to update a company's hybrid network to support IPv6 for the upcoming release of a new application. The application is hosted in a VPC in the AWS Cloud. The company's current AWS infrastructure includes VPCs that are connected by a transit gateway. The transit gateway is connected to the on-premises network by AWS Direct Connect and AWS Site-to-Site VPN. The company's on-premises devices have been updated to support the new IPv6 requirements.

The company has enabled IPv6 for the existing VPC by assigning a new IPv6 CIDR block to the VPC and by assigning IPv6 to the subnets for dual-stack support. The company has launched new Amazon EC2 instances for the new application in the updated subnets.

When updating the hybrid network to support IPv6 the network engineer must avoid making any changes to the current infrastructure. The network engineer also must block direct access to the instances' new IPv6 addresses from the internet. However, the network engineer must allow outbound internet access from the instances.

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

- A. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address
- B. Create a new VPN connection that supports IPv6 connectivity
- C. Add an egress-only internet gateway
- D. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices
- E. Update the Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address

- F. Update the existing VPN connection to support IPv6 connectivity
- G. Add an egress-only internet gateway
- H. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
- I. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address
- J. Create a new VPN connection that supports IPv6 connectivity
- K. Add an egress-only internet gateway
- L. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.
- M. Create a Direct Connect transit VIF and configure BGP peering with the AWS assigned IPv6 peering address
- N. Create a new VPN connection that supports IPv6 connectivity
- O. Add a NAT gateway
- P. Update any affected VPC security groups and route tables to provide connectivity within the VPC and between the VPC and the on-premises devices.

Answer: B

NEW QUESTION 21

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 22

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 27

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

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

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

Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 30

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 NLB. Use 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 ALB
- G. Use Bring Your Own IP (BYOIP) from the on-premises network with Global Accelerator
- H. Place the EC2 instances behind an Amazon CloudFront distribution
- I. Use Bring Your Own IP (BYOIP) from the on-premises network with CloudFront.

Answer: B

NEW QUESTION 34

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

- A. Create a new Site-to-Site VPN connection
- B. Set the transit gateway as the target gateway
- C. Enable acceleration on the new Site-to-Site VPN connection
- D. Update the VPN device in the London office with the new connection details.
- E. Modify the existing Site-to-Site VPN connection by setting the transit gateway as the target gateway. Enable acceleration on the existing Site-to-Site VPN connection.
- F. Create a new transit gateway in the eu-west-2 (London) Region
- G. Peer the new transit gateway with the existing transit gateway
- H. Modify the existing Site-to-Site VPN connection by setting the new transit gateway as the target gateway.
- I. Create a new AWS Global Accelerator standard accelerator that has an endpoint of the Site-to-Site VPN connection
- J. Update the VPN device in the London office with the new connection details.

Answer: A

Explanation:

Enabling acceleration for a Site-to-Site VPN connection uses AWS Global Accelerator to route traffic from the on-premises network to an AWS edge location that is closest to the customer gateway device¹. AWS Global Accelerator optimizes the network path, using the congestion-free AWS global network to route traffic to the endpoint that provides the best application performance². Setting the transit gateway as the target gateway enables connectivity between the on-premises network and multiple VPCs that are attached to the transit gateway³.

NEW QUESTION 35

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=true`
- F. Associate the new DHCP options set with the VPC.

Answer: B

NEW QUESTION 36

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

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

AWSTemplateFormatVersion: 2010-09-09 Parameters:

OriginatingVPCId: Type: String RemoteVPCId: Type: String

RemoteVPCAccountId: Type: String Resources:

newVPCPeeringConnection:

Type: 'AWS::EC2::VPCPeeringConnection'

Properties:

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

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

- A. Resources: NewEC2SecurityGroup: Type: AWS::EC2::SecurityGroup
- B. Resources: NetworkInterfaceToRemoteVPC: Type: "AWS::EC2::NetworkInterface"
- C. Resources: newEC2Route: Type: AWS::EC2::Route

D. Resources:VPCGatewayToRemoteVPC:Type: "AWS::EC2::VPCGatewayAttachment"
E. Resources:newVPCPeeringConnection:Type: 'AWS::EC2VPCPeeringConnection'PeerRoleArn: !Ref PeerRoleArn

Answer: CE

Explanation:

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

NEW QUESTION 40

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

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

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

Answer: A

NEW QUESTION 43

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 44

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

What should the network engineer do to meet this requirement?

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

Answer: D

NEW QUESTION 49

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) 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 50

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

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

What should the network engineer do to meet these requirements?

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

Answer: B

Explanation:

"If you advertise more than 100 routes each for IPv4 and IPv6 over the BGP session, the BGP session will go into an idle state with the BGP session DOWN." <https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html>

NEW QUESTION 51

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 54

A company has deployed a critical application on a fleet of Amazon EC2 instances behind an Application Load Balancer. The application must always be reachable on port 443 from the public internet. The application recently had an outage that resulted from an incorrect change to the EC2 security group.

A network engineer needs to automate a way to verify the network connectivity between the public internet and the EC2 instances whenever a change is made to the security group. The solution also must notify the network engineer when the change affects the connection.

Which solution will meet these requirements?

- A. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture REJECT traffic on port 443. Publish the flow log records to a log group in Amazon CloudWatch Log
- B. Create a CloudWatch Logs metric filter for the log group for rejected traffic
- C. Create an alarm to notify the network engineer.
- D. Enable VPC Flow Logs on the elastic network interface of each EC2 instance to capture all traffic on port 443. Publish the flow log records to a log group in Amazon CloudWatch Log
- E. Create a CloudWatch Logs metric filter for the log group for all traffic
- F. Create an alarm to notify the network engineer
- G. Create a VPC Reachability Analyzer path on port 443. Specify the security group as the source
- H. Specify the EC2 instances as the destination
- I. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection
- J. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs.
- K. Create a VPC Reachability Analyzer path on port 443. Specify the internet gateway of the VPC as the source
- L. Specify the EC2 instances as the destination
- M. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connection
- N. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fail
- O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function when a change to the security group occurs.

Answer: C

NEW QUESTION 56

A company has a global network and is using transit gateways to connect AWS Regions together. The company finds that two Amazon EC2 instances in different Regions are unable to communicate with each other. A network engineer needs to troubleshoot this connectivity issue. What should the network engineer do to meet this requirement?

- A. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables and in the VPC route table
- B. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.
- C. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables. Verify that the VPC route tables are correct
- D. Use AWS Firewall Manager to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.
- E. Use AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables. Verify that the VPC route tables are correct
- F. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.
- G. Use VPC Reachability Analyzer to analyze routes in the transit gateway route table
- H. Verify that the VPC route tables are correct
- I. Use VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC.

Answer: C

Explanation:

Using AWS Network Manager Route Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between VPCs and transit gateways¹. Verifying that the VPC route tables are correct would enable identification of routing issues within a VPC. Using VPC flow logs to analyze the IP traffic that security group rules and network ACL rules accept or reject in the VPC would enable identification of traffic filtering issues within a VPC². Additionally, using VPC Reachability Analyzer to analyze routes in the transit gateway route tables would enable identification of routing issues between transit gateways in different Regions. VPC Reachability Analyzer is a configuration analysis tool that enables connectivity testing between a source resource and a destination resource in your VPCs.

NEW QUESTION 59

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

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

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

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

Answer: C

NEW QUESTION 64

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 Amazon IP Space Changed 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

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 the existing subnets.

Answer: B

NEW QUESTION 66

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

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

Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 68

A company is deploying a new application in the AWS Cloud. The company wants a highly available web server that will sit behind an Elastic Load Balancer. The load balancer will route requests to multiple target groups based on the URL in the request. All traffic must use HTTPS. TLS processing must be offloaded to the load balancer. The web server must know the user's IP address so that the company can keep accurate logs for security purposes.

Which solution will meet these requirements?

- A. Deploy an Application Load Balancer with an HTTPS listener
- B. Use path-based routing rules to forward the traffic to the correct target group
- C. Include the X-Forwarded-For request header with traffic to the targets.
- D. Deploy an Application Load Balancer with an HTTPS listener for each domain
- E. Use host-based routing rules to forward the traffic to the correct target group for each domain
- F. Include the X-Forwarded-For request header with traffic to the targets.
- G. Deploy a Network Load Balancer with a TLS listener
- H. Use path-based routing rules to forward the traffic to the correct target group
- I. Configure client IP address preservation for traffic to the targets.
- J. Deploy a Network Load Balancer with a TLS listener for each domain
- K. Use host-based routing rules to forward the traffic to the correct target group for each domain
- L. Configure client IP address preservation for traffic to the targets.

Answer: A

Explanation:

An Application Load Balancer (ALB) can be used to route traffic to multiple target groups based on the URL in the request. The ALB can be configured with an HTTPS listener to ensure all traffic uses HTTPS. TLS processing can be offloaded to the ALB, which reduces the load on the web server. Path-based routing rules can be used to route traffic to the correct target group based on the URL in the request. The X-Forwarded-For request header can be included with traffic to the targets, which will allow the web server to know the user's IP address and keep accurate logs for security purposes.

NEW QUESTION 71

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

Which solution will meet these requirements without causing any downtime?

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

Answer: C

Explanation:

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

NEW QUESTION 73

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 77

A company has deployed a new web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Amazon EC2 Auto Scaling group. Enterprise customers from around the world will use the application. Employees of these enterprise customers will connect to the application over HTTPS from office locations.

The company must configure firewalls to allow outbound traffic to only approved IP addresses. The employees of the enterprise customers must be able to access the application with the least amount of latency.

Which change should a network engineer make in the infrastructure to meet these requirements?

- A. Create a new Network Load Balancer (NLB). Add the ALB as a target of the NLB.
- B. Create a new Amazon CloudFront distributio
- C. Set the ALB as the distribution's origin.
- D. Create a new accelerator in AWS Global Accelerato
- E. Add the ALB as an accelerator endpoint.
- F. Create a new Amazon Route 53 hosted zon
- G. Create a new record to route traffic to the ALB.

Answer: B

Explanation:

Amazon CloudFront is a content delivery network (CDN) that can speed up the delivery of static and dynamic web content, such as images, videos, and APIs². CloudFront can also provide end-to-end encryption for HTTPS traffic by using SSL certificates from AWS Certificate Manager (ACM) or other sources². CloudFront can also support session affinity (sticky sessions) with a load balancer-generated cookie or an application-based cookie policy².

NEW QUESTION 79

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:

AWS Global Accelerator can provide static IP addresses that the IoT devices can connect to without using DNS². It can also route traffic over the AWS global network and improve performance and availability for the IoT devices². An NLB can provide end-to-end encryption for HTTPS traffic by using TLS as a target group protocol and terminating SSL connections at the load balancer level¹. An NLB can also support session affinity (sticky sessions) with TCP connections¹.

NEW QUESTION 83

A company wants to improve visibility into its AWS environment. The AWS environment consists of multiple VPCs that are connected to a transit gateway. The transit gateway connects to an on-premises data center through an AWS Direct Connect gateway and a pair of redundant Direct Connect connections that use

transit VIFs. The company must receive notification each time a new route is advertised to AWS from on premises over Direct Connect. What should a network engineer do to meet these requirements?

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

Answer: B

Explanation:

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

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

NEW QUESTION 88

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

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

Answer: A

NEW QUESTION 90

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

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

What should a network engineer do to meet these requirements?

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 92

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 forma
- 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 forma
- 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 forma
- 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 forma
- 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 97

A company is hosting an application on Amazon EC2 instances behind a Network Load Balancer (NLB). A solutions architect added EC2 instances in a second Availability Zone to improve the availability of the application. The solutions architect added the instances to the NLB target group. The company's operations team notices that traffic is being routed only to the instances in the first Availability Zone. What is the MOST operationally efficient solution to resolve this issue?

- A. Enable the new Availability Zone on the NLB
- B. Create a new NLB for the instances in the second Availability Zone
- C. Enable proxy protocol on the NLB
- D. Create a new target group with the instances in both Availability Zones

Answer: A

Explanation:

When adding instances in a new Availability Zone to an existing Network Load Balancer (NLB), it is important to ensure that the new Availability Zone is enabled on the NLB. This will allow traffic to be routed to instances in both Availability Zones. This can be done by editing the settings of the NLB and selecting the new Availability Zone from the list of available zones.

NEW QUESTION 100

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 VPN to connect the virtual private gateways to the corporate network.
- E. Create a Direct Connect gateway, and add virtual private gateway associations to the VPC
- F. Configure a private VIF to connect to the corporate network.
- G. Create a transit gateway, and attach the VPC
- H. Create a Direct Connect gateway, and associate it with the transit 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 transit gateway, and on-premises networks can communicate with the VPCs through the Direct Connect gateway. This solution provides a central point of management and simplifies the configuration of network routing. By associating the Direct Connect gateway with the transit gateway, traffic between the VPCs and the on-premises network can be routed through the Direct Connect connection.

NEW QUESTION 103

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 105

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

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

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

- A. Enable VPC flow logs on the NAT gateway's elastic network interfac
- B. Publish the logs to a log group in Amazon CloudWatch Log
- C. Use CloudWatch Logs Insights to query and analyze the logs.
- D. Enable NAT gateway access log
- E. Publish the logs to a log group in Amazon CloudWatch Log
- F. Use CloudWatch Logs Insights to query and analyze the logs.
- G. Configure Traffic Mirroring on the NAT gateway's elastic network interfac
- H. Send the traffic to an additional EC2 instanc
- I. Use tools such as tcpdump and Wireshark to query and analyze the mirrored traffic.
- J. Enable VPC flow logs on the NAT gateway's elastic network interfac

- K. Publish the logs to an Amazon S3 bucket
- L. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure
- M. Use Athena to query and analyze the logs.
- N. Enable NAT gateway access log
- O. Publish the logs to an Amazon S3 bucket
- P. Create a custom table for the S3 bucket in Amazon Athena to describe the log structure
- Q. Use Athena to query and analyze the logs.

Answer: AD

Explanation:

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

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

(Option A)

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

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

NEW QUESTION 107

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

The SQS queue is not receiving messages.

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

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

Answer: CE

NEW QUESTION 111

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

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

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

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

Answer: BC

NEW QUESTION 113

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 117

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

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

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 119

.....

Relate Links

100% Pass Your ANS-C01 Exam with ExamBible Prep Materials

<https://www.exambible.com/ANS-C01-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/>