# ANS-C01 Dumps

# AWS Certified Advanced Networking Specialty Exam

## https://www.certleader.com/ANS-C01-dumps.html

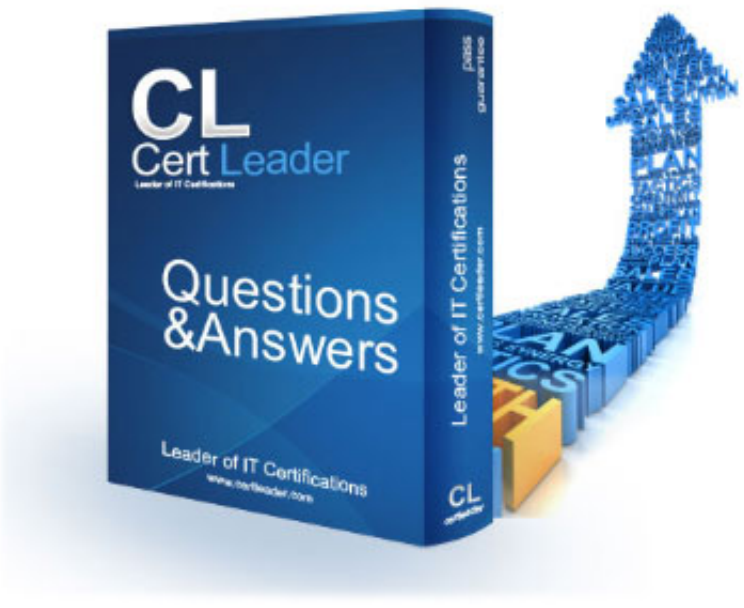

**NEW QUESTION 1**
A company has an AWS Site-to-Site VPN connection between its existing VPC and on-premises network. The default DHCP options set is associated with the VPC. The company has an application that is running on an Amazon Linux 2 Amazon EC2 instance in the VPC. The application must retrieve 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 VP
C. Reboot the Amazon Linux 2 EC2 instance.
D. Create an Amazon Route 53 Resolver rul
E. Associate the rule with the VP
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 VPMap 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 VP
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)3. 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 endpoints2.

**NEW QUESTION 2**
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 endpoin
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 endpoin
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 endpoin
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 endpoin
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 3**
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 Gatewa
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 VP
E. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
F. Ensure that private DNS is turned of
G. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
H. Create an Amazon Route 53 forwarding rule for each interface VPC endpoin

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 VPIn the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
L. Ensure that private DNS is turned of
M. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
N. Create an Amazon Route 53 private hosted zone with a full service endpoint name for Amazon S3 and Systems Manage
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 VP
R. In the central shared services VPC, create interface VPC endpoints for Amazon S3 and Systems Manager to acces
S. Connect all the VPCs to the central shared services VPC by using AWS Transit Gatewa
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 serviceswithout requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection2. Interface VPC endpoints are powered by AWS PrivateLink, a technology that enables private access to AWS services2. Amazon S3 and AWS Systems Manager support interface VPC endpoin2ts. By turning off private DNS, the interface VPC endpoints can be accessed by using their private IP addresses2. By using Amazon Route 53 forwarding rules, DNS queries can be resolved to the interface VPC endpoints in the shared services VPC3.

**NEW QUESTION 4**
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 subne
B. Disable the Auto-assign Public IP option while launching the EC2 instance
C. Create an internet gatewa
D. Attach the internet gateway to the VP
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 subne
G. Create a NAT gateway in a public subnet in the same Availability Zon
H. Create an internet gatewa
I. Attach the internet gateway to the VP
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 subne
L. Create an interface VPC endpoint for Amazon SQ
M. Create gateway VPC endpoints for Amazon S3 and DynamoDB.
N. Place the EC2 instances in a private subne
O. Create a gateway VPC endpoint for Amazon SQS.Create interface VPC endpoints for Amazon S3 and DynamoDB.

**Answer:** C

**NEW QUESTION 5**
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 6**
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 7**
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 8**
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 9**
A customer has set up multiple VPCs for Dev, Test, Prod, and Management. You need to set up AWS Direct Connect to enable data flow from on-premises to each VPC. The customer hasmonitoring software running in the Management VPC that collects metrics from the instances in all the other VPCs. Due to budget requirements, data transfer charges should be kept at minimum.
Which design should be recommended?

A. Create a total of four private VIFs, one for each VPC owned by the customer, and route traffic between VPCs using the Direct Connect link.
B. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs.
C. Create a private VIF to the Management VPC, and peer this VPC to all other VPCs, enable source/destination NAT in the Management VPC.

D. Create a total of four private VIFs, and enable VPC peering between all VPCs.

**Answer:** D

**Explanation:**
- creating VPC peering is free of charge - traffic costs ~0.01€/GB for VPC peering (IN + OUT) and ~0.02€/GB for direct connect (OUT only). As the communication involved in monitoring will never have IN == OUT, then 0.01 * (IN + OUT) will always be lower the 0.02 * OUT, ergo VPC peering will be cheaper

**NEW QUESTION 10**
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 addres
B. Create a new VPN connection that supports IPv6 connectivit
C. Add an egress-only internet gatewa
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 addres
F. Update the existing VPN connection to support IPv6 connectivit
G. Add an egress-only internet gatewa
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 addres
J. Create a new VPN connection that supports IPv6 connectivit
K. Add an egress-only internet gatewa
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 addres
N. Create a new VPN connection that supports IPv6 connectivit
O. Add a NAT gatewa
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 10**
You deploy an Amazon EC2 instance that runs a web server into a subnet in a VPC. An Internet gateway is attached, and the main route table has a default route (0.0.0.0/0) configured with a target of the Internet gateway.
The instance has a security group configured to allow as follows:
- Protocol: TCP
- Port: 80 inbound, nothing outbound

The Network ACL for the subnet is configured to allow as follows:
- Protocol: TCP
- Port: 80 inbound, nothing outbound

When you try to browse to the web server, you receive no response. Which additional step should you take to receive a successful response?

A. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 80
B. Add an entry to the security group outbound rules for Protocol: TCP, Port Range: 1024-65535
C. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 80
D. Add an entry to the Network ACL outbound rules for Protocol: TCP, Port Range: 1024-65535

**Answer:** D

**Explanation:**
To enable the connection to a service running on an instance, the associated network ACL must allow both inbound traffic on the port that the service is listening on as well as allow outbound traffic from ephemeral ports. When a client connects to a server, a random port from the ephemeral port range (1024-65535) becomes the client's source port. The designated ephemeral port then becomes the destination port for return traffic from the service, so outbound traffic from the ephemeral port must be allowed in the network ACL.https://aws.amazon.com/premiumsupport/knowledge-center/resolve-connection-sg-acl-inbound/

**NEW QUESTION 11**
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 applicatio
B. Create a link aggregation group (LAG).
C. Deploy an AWS Site-to-Site VPN connection to the application VP
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 VPInstruct the remote employees to connect to Workspaces.
F. Replace the existing 1 Gbps Direct Connect connection with two new 2 Gbps Direct Connect hosted connection
G. Create an AWS Client VPN endpoint in the application VP
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 trafficload 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 15**
A company is building its website on AWS in a single VPC. The VPC has public subnets and private subnets in two Availability Zones. The website has static content such as images. The company is using Amazon S3 to store the content.
The company has deployed a fleet of Amazon EC2 instances as web servers in a private subnet. The EC2 instances are in an Auto Scaling group behind an Application Load Balancer. The EC2 instances will serve traffic, and they must pull content from an S3 bucket to render the webpages. The company is using AWS Direct Connect with a public VIF for on-premises connectivity to the S3 bucket.
A network engineer notices that traffic between the EC2 instances and Amazon S3 is routing through a NAT gateway. As traffic increases, the company's costs are increasing. The network engineer needs to change the connectivity to reduce the NAT gateway costs that result from the traffic between the EC2 instances and Amazon S3.
Which solution will meet these requirements?

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

**Answer:** D

**NEW QUESTION 17**
An IoT company sells hardware sensor modules that periodically send out temperature, humidity, pressure, and location data through the MQTT messaging protocol. The hardware sensor modules send this data to the company's on-premises MQTT brokers that run on Linux servers behind a load balancer. The hardware sensor modules have been hardcoded with public IP addresses to reach the brokers.
The company is growing and is acquiring customers across the world. The existing solution can no longer scale and is introducing additional latency because of the company's global presence. As a result, the company decides to migrate its entire infrastructure from on premises to the AWS Cloud. The company needs to migrate without reconfiguring the hardware sensor modules that are already deployed across the world. The solution also must minimize latency.
The company migrates the MQTT brokers to run on Amazon EC2 instances. What should the company do next to meet these requirements?

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

**Answer:** B

**NEW QUESTION 21**
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 tom 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 connectio
B. Set the transit gateway as the target gatewa
C. Enable acceleration on the new Site-to-Site VPN connectio
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) Regio
G. Peer the new transit gateway with the existing transit gatewa
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 connectio
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 26**
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 VP
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 AWSaccount to resolve queries for this domain.
G. Launch two Amazon EC2 instances in the shared AWS accoun
H. Install BIND on each instanc
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 accoun
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 30**
A company has deployed a software-defined WAN (SD-WAN) solution to interconnect all of its offices. The company is migrating workloads to AWS and needs to extend its SD-WAN solution to support connectivity to these workloads.
A network engineer plans to deploy AWS Transit Gateway Connect and two SD-WAN virtual appliances to provide this connectivity. According to company policies, only a single SD-WAN virtual appliance can handle traffic from AWS workloads at a given time.
How should the network engineer configure routing to meet these requirements?

A. Add a static default route in the transit gateway route table to point to the secondary SD-WAN virtual applianc
B. Add routes that are more specific to point to the primary SD-WAN virtual appliance.
C. Configure the BGP community tag 7224:7300 on the primary SD-WAN virtual appliance for BGP routes toward the transit gateway.
D. Configure the AS_PATH prepend attribute on the secondary SD-WAN virtual appliance for BGP routes toward the transit gateway.
E. Disable equal-cost multi-path (ECMP) routing on the transit gateway for Transit Gateway Connect.

**Answer:** A

**NEW QUESTION 33**
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:
AWSTemplateFormation Version: 2010-09-09 Parameters:
Originating VCId: Type: String RemoteVPCId: Type: String
RemoteVPCAccountId: Type: String Resources:
newVPCPeeringConnection:
Type: 'AWS::EC2::VPCPeeringConnection'
Properties:
VpcdId: !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::EC2NetworkInterface"
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 36**
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 servic
B. Create a lifecycle hook to add new tasks to the target group from Amazon ECS as required to handle scalin
C. Specify the GLB in the service definitio
D. Create a VPC peer for external AWS account
E. Update the route tables so that the AWS accounts can reach the GLB.
F. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS servic
G. Create path-based routing rules to allow the application to target the containers that are registered in the target grou
H. Specify the ALB in the service definitio
I. Create a VPC endpoint service for the ALB Share the VPC endpoint service with other AWS accounts.
J. Choose an Application Load Balancer (ALB) as the type of load balancer for the ECS servic
K. Create path-based routing rules to allow the application to target the containers that are registered in the target grou
L. Specify the ALB in the service definitio
M. Create a VPC peer for the external AWS account
N. Update the route tables so that the AWS accounts can reach the ALB.
O. Choose a Network Load Balancer (NLB) as the type of load balancer for the ECS servic
P. Specify the NLB in the service definitio
Q. Create a VPC endpoint service for the NL
R. Share the VPC endpoint service with other AWS accounts.

**Answer:** D

**NEW QUESTION 39**
A company plans to deploy a two-tier web application to a new VPC in a single AWS Region. The company has configured the VPC with an internet gateway and four subnets. Two of the subnets are public and have default routes that point to the internet gateway. Two of the subnets are private and share a route table that does not have a default route.
The application will run on a set of Amazon EC2 instances that will be deployed behind an external Application Load Balancer. The EC2 instances must not be directly accessible from the internet. The application will use an Amazon S3 bucket in the same Region to store data. The application will invoke S3 GET API operations and S3 PUT API operations from the EC2 instances. A network engineer must design a VPC architecture that minimizes data transfer cost.
Which solution will meet these requirements?

A. Deploy the EC2 instances in the public subnet
B. Create an S3 interface endpoint in the VP
C. Modify the application configuration to use the S3 endpoint-specific DNS hostname.
D. Deploy the EC2 instances in the private subnet
E. Create a NAT gateway in the VP
F. Create default routes in the private subnets to the NAT gatewa
G. Connect to Amazon S3 by using the NAT gateway.
H. Deploy the EC2 instances in the private subnet
I. Create an S3 gateway endpoint in the VPSpecify die route table of the private subnets during endpoint creation to create routes to Amazon S3.
J. Deploy the EC2 instances in the private subnet
K. Create an S3 interface endpoint in the VP
L. Modify the application configuration to use the S3 endpoint-specific DNS hostname.

**Answer:** C

**Explanation:**
Option C is the optimal solution as it involves deploying the EC2 instances in the private subnets, which provides additional security benefits. Additionally, creating an S3 gateway endpoint in the VPC will enable the EC2 instances to communicate with Amazon S3 directly, without incurring data transfer costs. This is because the S3 gateway endpoint uses Amazon's private network to transfer data between the VPC and S3, which is not charged for data transfer. Furthermore, specifying the route table of the private subnets during endpoint creation will create routes to Amazon S3, which is required for the EC2 instances to communicate with S3.

**NEW QUESTION 40**
A company is using a NAT gateway to allow internet connectivity for private subnets in a VPC in the us-west-2 Region. After a security audit, the company needs to remove the NAT gateway.
In the private subnets, the company has resources that use the unified Amazon CloudWatch agent. A network engineer must create a solution to ensure that the unified CloudWatch agent continues to work after the removal of the NAT gateway.
Which combination of steps should the network engineer take to meet these requirements? (Choose three.)

A. Validate that private DNS is enabled on the VPC by setting the enableDnsHostnames VPC attribute and the enableDnsSupport VPC attribute to true.
B. Create a new security group with an entry to allow outbound traffic that uses the TCP protocol on port 443 to destination 0.0.0.0/0
C. Create a new security group with entries to allow inbound traffic that uses the TCP protocol on port 443 from the IP prefixes of the private subnets.
D. Create the following interface VPC endpoints in the VPC: com.amazonaws.us-west-2.logs and com.amazonaws.us-west-2.monitorin
E. Associate the new security group with the endpoint network interfaces.
F. Create the following interface VPC endpoint in the VPC: com.amazonaws.us-west-2.cloudwatch.Associate the new security group with the endpoint network interfaces.
G. Associate the VPC endpoint or endpoints with route tables that the private subnets use.

**Answer:** BDF

**NEW QUESTION 45**
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-proto

**NEW QUESTION 46**
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 47**
A company has deployed an AWS Network Firewall firewall into a VPC. A network engineer needs to implement a solution to deliver Network Firewall flow logs to the company's Amazon OpenSearch Service (Amazon Elasticsearch Service) cluster in the shortest possible time.
Which solution will meet these requirements?

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

**Answer:** B

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

**NEW QUESTION 50**
A global company operates all its non-production environments out of three AWS Regions: eu-west-1,
us-east-1, and us-west-1. The company hosts all its production workloads in two on-premises data centers. The company has 60 AWS accounts and each account has two VPCs in each Region. Each VPC has a virtual private gateway where two VPN connections terminate for resilient connectivity to the data centers. The company has 360 VPN tunnels to each data center, resulting in high management overhead. The total VPN throughput for each Region is 500 Mbps.
The company wants to migrate the production environments to AWS. The company needs a solution that will simplify the network architecture and allow for future growth. The production environments will generate an additional 2 Gbps of traffic per Region back to the data centers. This traffic will increase over time.
Which solution will meet these requirements?

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

**Answer:** C

**NEW QUESTION 54**
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 Gatewa
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 57**
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 traffi
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 traffi
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 sourc
H. Specify the EC2 instances as the destinatio
I. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connectio
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 groupoccurs.
K. Create a VPC Reachability Analyzer path on port 443. Specify the internet gateway of the VPC as the sourc
L. Specify the EC2 instances as the destinatio
M. Create an Amazon Simple Notification Service (Amazon SNS) topic to notify the network engineer when a change to the security group affects the connectio
N. Create an AWS Lambda function to start Reachability Analyzer and to publish a message to the SNS topic in case the analyses fai
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 58**
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 correc
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 correc
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 correc
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**
An organization launched an IPv6-only web portal to support IPv6-native mobile clients. Front-end instances launch in an Amazon VPC associated with an appropriate IPv6 CIDR. The VPC IPv4 CIDR is fully utilized. A single subnet exists in each of two Availability Zones with appropriately configured IPv6 CIDR associations. Auto Scaling is properly configured, and no Elastic Load Balancing is used.
Customers say the service is unavailable during peak load times. The network engineer attempts to launch an instance manually and receives the following message: “There are not enough free addresses in subnet ‘subnet-12345677’ to satisfy the requested number of instances.”
What action will resolve the availability problem?

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

**Answer:** B

**NEW QUESTION 60**
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 applianc
B. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gatewa
C. Configure BGP over the IPsec VPN connections
D. Assign a new CIDR block to the transit gatewa
E. Create a new VPC for the SD-WAN hub virtual applianc
F. Attach the new VPC to the transit gateway with a VPC attachmen
G. Add a transit gateway Connect attachmen
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 applianc
K. Attach the new VPC to the transit gateway with a VPC attachmen
L. Create two IPsec VPN connections between the SD-WAN hub virtual appliance and the transit gatewa
M. Configure BGP over the IPsec VPN connections.
N. Assign a new CIDR block to the transit gatewa
O. Create a new VPC for the SD-WAN hub virtual applianc
P. Attach the new VPC to the transit gateway with a VPC attachmen
Q. Add a transit gateway Connect attachmen
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 64**
A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is the origin in an Amazon CloudFront distribution. The company wants to implement a custom authentication system that will provide a token for its authenticated customers.
The web application must ensure that the GET/POST requests come from authenticated customers before it delivers the content. A network engineer must design a solution that gives the web application the ability to identify authorized customers.
What is the MOST operationally efficient solution that meets these requirements?

A. Use the ALB to inspect the authorized token inside the GET/POST request payloa
B. Use an AWS Lambda function to insert a customized header to inform the web application of an authenticated customer request.
C. Integrate AWS WAF with the ALB to inspect the authorized token inside the GET/POST request payloa
D. Configure the ALB listener to insert a customized header to inform the web application of an authenticated customer request.
E. Use an AWS Lambda@Edge function to inspect the authorized token inside the GET/POST request payloa
F. Use the Lambda@Edge function also to insert a customized header to inform the web application of an authenticated customer request.
G. Set up an EC2 instance that has a third-party packet inspection tool to inspect the authorized token inside the GET/POST request payloa
H. Configure the tool to insert a customized header to inform the web application of an authenticated customer request.

**Answer:** C

**NEW QUESTION 65**
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 listene
B. Use path-based routing rules to forward the traffic to the correct target grou
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 domai
E. Use host-based routing rules to forward the traffic to the correct target group for each domai
F. Include the X-Forwarded-For request header with traffic to the targets.
G. Deploy a Network Load Balancer with a TLS listene
H. Use path-based routing rules to forward the traffic to the correct target grou
I. Configure client IP address preservation for traffic to the targets.
J. Deploy a Network Load Balancer with a TLS listener for each domai
K. Use host-based routing rules to forward the traffic to the correct target group for each domai
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 66**
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 70**
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 packag
B. Wait until the analysis has finished running to find out which EC2 instances are still listening to the old port.
C. Enable Amazon GuardDut
D. Use the graphical visualizations to filter for traffic that uses the port of the old protoco
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 bucke
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 port
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 VPC3. 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 75**
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) endpoin
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) endpoin
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 76**
A bank built a new version of its banking application in AWS using containers that content 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 applicatio
C. Route all traffic to the new application by using an Elastic Load Balancing (ELB) load balancer DN
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 applicatio
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 applicatio
H. Register both the new and earlier application backends as separate target group
I. Use header-based routing to route traffic based on the application version.

**Answer:** D

**NEW QUESTION 78**
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 83**
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 85**
A company is migrating an application from on premises to AWS. The company will host the application on Amazon EC2 instances that are deployed in a single VPC. During the migration period, DNS queries from the EC2 instances must be able to resolve names of on-premises servers. The migration is expected to take 3 months After the 3-month migration period, the resolution of on-premises servers will no longer be needed.
What should a network engineer do to meet these requirements with the LEAST amount of configuration?

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

**Answer:** A

**Explanation:**
Setting up an AWS Site-to-Site VPN connection between on premises and AWS would enable a secure and encrypted connection over the public internet1. Deploying an Amazon Route 53 Resolver outbound endpoint in the Region that is hosting the VPC would enable forwarding of DNS queries for on-premises servers to the on-premises DNS servers2. This would allow EC2 instances in the VPC to resolve names of on-premises servers during the migration period. After the migration period, the Route 53 Resolver outbound endpoint can be deleted with minimal configuration changes.

**NEW QUESTION 87**
A company uses AWS Direct Connect to connect its corporate network to multiple VPCs in the same AWS account and the same AWS Region. Each VPC uses its own private VIF and its own virtual LAN on the Direct Connect connection. The company has grown and will soon surpass the limit of VPCs and private VIFs for each connection.
What is the MOST scalable way to add VPCs with on-premises connectivity?

A. Provision a new Direct Connect connection to handle the additional VPC
B. Use the new connection to connect additional VPCs.
C. Create virtual private gateways for each VPC that is over the service quot
D. Use AWS Site-to-Site VPNto connect the virtual private gateways to the corporate network.
E. Create a Direct Connect gateway, and add virtual private gateway associations to the VPC
F. Configure a private VIF to connect to the corporate network.
G. Create a transit gateway, and attach the VPC
H. Create a Direct Connect gateway, and associate it with the transit gatewa
I. Create a transit VIF to the Direct Connect gateway.

**Answer:** D

**Explanation:**
When a company requires connectivity to multiple VPCs over AWS Direct Connect, a scalable solution is to use a transit gateway. A transit gateway is a hub that can interconnect multiple VPCs and VPN connections. The VPCs can communicate with each other over the transitgateway, and on-premises networks can communicate with the VPCs through the Direct Connect gateway. This solution provides a central point of management and simplifies the configuration of network routing. By associating the Direct Connect gateway with the transit gateway, traffic between the VPCs and the on-premises network can be routed through the Direct Connect connection.

**NEW QUESTION 88**
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 92**
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 VP
B. Connect the inspection VPC to the transit gateway by using a VPCattachmen
C. Create a target group, and register the appliances with the target grou
D. Create a Network Load Balancer (NLB), and set it up to forward to the newly created target grou
E. Configure a default route in the inspection VPCs transit gateway subnet toward the NLB.
F. Deploy two clusters that consist of multiple appliances across multiple Availability Zones in a designated inspection VP
G. Connect the inspection VPC to the transit gateway by using a VPC attachmen
H. Create a target group, and register the appliances with the target grou
I. Create a Gateway Load Balancer, and set it up to forward to the newly created target grou
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 gatewa

L. Associate one route table with all the attachments of the application VPC
M. Associate the other route table with the inspection VPC's attachmen
N. Propagate all VPC attachments into the inspection route tabl
O. Define a static default route in the application route tabl
P. Enable appliance mode on the attachment that connects the inspection VPC.
Q. Configure two route tables on the transit gatewa
R. Associate one route table with all the attachments of the application VPC
S. Associate the other route table with the inspection VPCs attachmen
T. Propagate all VPC attachments into the application route tabl
. Define a static default route in the inspection route tabl
. Enable appliance mode on the attachment that connects the inspection VPC.
. Configure one route table on the transit gatewa
. Associate the route table with all the VPC
. Propagate all VPC attachments into the route tabl
. Define a static default route in the route table.

**Answer:** BC

**NEW QUESTION 93**
.....

# Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

## 100% Pass Your ANS-C01 Exam with Our Prep Materials Via below:

https://www.certleader.com/ANS-C01-dumps.html