# Exam Questions ANS-C01

AWS Certified Advanced Networking Specialty Exam

## https://www.2passeasy.com/dumps/ANS-C01/

**NEW QUESTION 1**
A company is using custom DNS servers that run BIND for name resolution in its VPCs. The VPCs are deployed across multiple AWS accounts that are part of the same organization in AWS Organizations. All the VPCs are connected to a transit gateway. The BIND servers are running in a central VPC and are configured to forward all queries for an on-premises DNS domain to DNS servers that are hosted in an on-premises data center. To ensure that all the VPCs use the custom DNS servers, a network engineer has configured a VPC DHCP options set in all the VPCs that specifies the custom DNS servers to be used as domain name servers.

Multiple development teams in the company want to use Amazon Elastic File System (Amazon EFS). A development team has created a new EFS file system but cannot mount the file system to one of its Amazon EC2 instances. The network engineer discovers that the EC2 instance cannot resolve the IP address for the EFS mount point fs-33444567d.efs.us-east-1.amazonaws.com. The network engineer needs to implement a solution so that development teams throughout the organization can mount EFS file systems.
Which combination of steps will meet these requirements? (Choose two.)

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

**Answer:** BD

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

**NEW QUESTION 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 is planning to use Amazon S3 to archive financial data. The data is currently stored in an
on-premises data center. The company uses AWS Direct Connect with a Direct Connect gateway and a transit gateway to connect to the on-premises data center.
The data cannot be transported over the public internet and must be encrypted in transit.
Which solution will meet these requirements?

A. Create a Direct Connect public VI
B. Set up an IPsec VPN connection over the public VIF to access Amazon S3. Use HTTPS for communication.
C. Create an IPsec VPN connection over the transit VI

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

**Answer:** B

**Explanation:**
https://docs.aws.amazon.com/vpn/latest/s2svpn/private-ip-dx.html
An IPsec VPN connection over the transit VIF can encrypt traffic between the on-premises network and AWS without using public IP addresses or the internet2. A VPC endpoint for Amazon S3 can enable private access to S3 buckets within the same region.HTTPS can provide additional encryption for communication.


**NEW QUESTION 4**
A company has an AWS Direct Connect connection between its on-premises data center in the United States (US) and workloads in the us-east-1 Region. The connection uses a transit VIF to connect the data center to a transit gateway in us-east-1.
The company is opening a new office in Europe with a new on-premises data center in England. A Direct Connect connection will connect the new data center with some workloads that are running in a single VPC in the eu-west-2 Region. The company needs to connect the US data center and us-east-1 with the Europe data center and eu-west-2. A network engineer must establish full connectivity between the data centers and Regions with the lowest possible latency.
How should the network engineer design the network architecture to meet these requirements?

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

**Answer:** C


**NEW QUESTION 5**
A company has deployed a web application on AWS. The web application uses an Application Load Balancer (ALB) across multiple Availability Zones. The targets of the ALB are AWS Lambda functions. The web application also uses Amazon CloudWatch metrics for monitoring.
Users report that parts of the web application are not loading properly. A network engineer needs to troubleshoot the problem. The network engineer enables access logging for the ALB.
What should the network engineer do next to determine which errors the ALB is receiving?

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

**Answer:** B

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


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

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

**Answer:** B

**NEW QUESTION 8**
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 i
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 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 company has two AWS accounts one for Production and one for Connectivity. A network engineer needs to connect the Production account VPC to a transit gateway in the Connectivity account. The feature to auto accept shared attachments is not enabled on the transit gateway.
Which set of steps should the network engineer follow in each AWS account to meet these requirements?

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

**Answer:** A

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

**NEW QUESTION 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 14**
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 18**
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 22**
A global delivery company is modernizing its fleet management system. The company has several business units. Each business unit designs and maintains applications that are hosted in its own AWS account in separate application VPCs in the same AWS Region. Each business unit's applications are designed to get data from a central shared services VPC.
The company wants the network connectivity architecture to provide granular security controls. The architecture also must be able to scale as more business units consume data from the central shared services VPC in the future.
Which solution will meet these requirements in the MOST secure manner?

A. Create a central transit gatewa
B. Create a VPC attachment to each application VP
C. Provide full mesh connectivity between all the VPCs by using the transit gateway.
D. Create VPC peering connections between the central shared services VPC and each application VPC in each business unit's AWS account.
E. Create VPC endpoint services powered by AWS PrivateLink in the central shared services VPCCreate VPC endpoints in each application VPC.
F. Create a central transit VPC with a VPN appliance from AWS Marketplac
G. Create a VPN attachment from each VPC to the transit VP
H. Provide full mesh connectivity among all the VPCs.

**Answer:** C

**Explanation:**
Option C provides a secure and scalable solution using VPC endpoint services powered by AWS PrivateLink. AWS PrivateLink enables private connectivity between VPCs and services without exposing the data to the public internet or using a VPN connection. By creating VPC endpoints in each application VPC, the company can securely access the central shared services VPC without the need for complex network configurations. Furthermore, PrivateLink supports cross-account connectivity, which makes it a scalable solution as more business units consume data from the central shared services VPC in the future.

**NEW QUESTION 24**
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 26**
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 31**
A company manages resources across VPCs in multiple AWS Regions. The company needs to connect to the resources by using its internal domain name. A network engineer needs to apply the aws.example.com DNS suffix to all resources.
What must the network engineer do to meet this requirement?

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

**Answer:** D

**Explanation:**
Creating one private hosted zone for aws.example.com and associating it with every VPC that has resources would enable DNS resolution for all resources by using their internal domain name. Creating an alias record in each private hosted zone with the full AWS service endpoint pointing to the interface VPC endpoint in the shared services VPC would enable private connectivity to Amazon S3 and AWS Systems Manager without using public endpoints.

**NEW QUESTION 34**
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 38**
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 39**
A company has a hybrid cloud environment. The company's data center is connected to the AWS Cloud by an AWS Direct Connect connection. The AWS environment includes VPCs that are connected together in a hub-and-spoke model by a transit gateway. The AWS environment has a transit VIF with a Direct Connect gateway for on-premises connectivity.
The company has a hybrid DNS model. The company has configured Amazon Route 53 Resolver endpoints in the hub VPC to allow bidirectional DNS traffic flow. The company is running a backend application in one of the VPCs.
The company uses a message-oriented architecture and employs Amazon Simple Queue Service (Amazon SQS) to receive messages from other applications over a private network. A network engineer wants to use an interface VPC endpoint for Amazon SQS for this architecture. Client services must be able to access the endpoint service from on premises and from multiple VPCs within the company's AWS infrastructure.
Which combination of steps should the network engineer take to ensure that the client applications can resolve DNS for the interface endpoint? (Choose three.)

A. Create the interface endpoint for Amazon SQS with the option for private DNS names turned on.
B. Create the interface endpoint for Amazon SQS with the option for private DNS names turned off.
C. Manually create a private hosted zone for sqs.us-east-1.amazonaws.co
D. Add necessary records that point to the interface endpoin
E. Associate the private hosted zones with other VPCs.
F. Use the automatically created private hosted zone for sqs.us-east-1.amazonaws.com with previously created necessary records that point to the interface endpoin
G. Associate the private hosted zones with other VPCs.
H. Access the SQS endpoint by using the public DNS name sqs.us-east-1 amazonaws.com in VPCs and on premises.
I. Access the SQS endpoint by using the private DNS name of the interface endpoint.sqs.us-east-1.vpce.amazonaws.com in VPCs and on premises.

**Answer:** ADF


**NEW QUESTION 40**
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) duster 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 Kubernete
B. Using that controller, configure a NetworkLoad 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 Kubernete
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 grou
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 grou
H. Add the EKS managed node group's Auto Scaling group as a targe
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-gro


**NEW QUESTION 43**
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 47**
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 48**
An international company provides early warning about tsunamis. The company plans to use IoT devices to monitor sea waves around the world. The data that is collected by the IoT devices must reach the company's infrastructure on AWS as quickly as possible. The company is using three operation centers around the world. Each operation center is connected to AWS through Its own AWS Direct Connect connection. Each operation center is connected to the internet through at least two upstream internet service providers.
The company has its own provider-independent (PI) address space. The IoT devices use TCP protocols for reliable transmission of the data they collect. The IoT devices have both landline and mobile internet connectivity. The infrastructure and the solution will be deployed in multiple AWS Regions. The company will use Amazon Route 53 for DNS services.
A network engineer needs to design connectivity between the IoT devices and the services that run in the AWS Cloud.
Which solution will meet these requirements with the HIGHEST availability?

A. Set up an Amazon CloudFront distribution with origin failove
B. Create an origin group for each Region where the solution is deployed.
C. Set up Route 53 latency-based routin
D. Add latency alias record
E. For the latency alias records, set the value of Evaluate Target Health to Yes.
F. Set up an accelerator in AWS Global Accelerato
G. Configure Regional endpoint groups andhealth checks.
H. Set up Bring Your Own IP (BYOIP) addresse
I. Use the same PI addresses for each Region where the solution is deployed.

**Answer:** B

**Explanation:**
https://aws.amazon.com/blogs/iot/automate-global-device-provisioning-with-aws-iot-core-and-amazon-route-53


**NEW QUESTION 52**
A company's network engineer builds and tests network designs for VPCs in a development account. The company needs to monitor the changes that are made to network resources and must ensure strict compliance with network security policies. The company also needs access to the historical configurations of network resources.
Which solution will meet these requirements?

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a custom pattern to monitor the account for change
B. Configure the rule to invoke an AWS Lambda function to identify noncompliant resource
C. Update an Amazon DynamoDB table with the changes that are identified.
D. Create custom metrics from Amazon CloudWatch log
E. Use the metrics to invoke an AWS Lambda function to identify noncompliant resource
F. Update an Amazon DynamoDB table with the changes that are identified.
G. Record the current state of network resources by using AWS Confi
H. Create rules that reflect the desired configuration setting
I. Set remediation for noncompliant resources.
J. Record the current state of network resources by using AWS Systems Manager Inventor
K. Use Systems Manager State Manager to enforce the desired configuration settings and to carry out remediation for noncompliant resources.

**Answer:** C

**Explanation:**
Recording the current state of network resources by using AWS Config would enable auditing and assessment of resource configurations and compliance3.
Creating rules that reflect the desired configuration settings would enable evaluation of whether the network resources comply with network security policies3.
Setting remediation for noncompliant resources would enable automatic correction of undesired configurations3.


**NEW QUESTION 55**
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 59**
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 61**
A company uses a 4 Gbps AWS Direct Connect dedicated connection with a link aggregation group (LAG) bundle to connect to five VPCs that are deployed in the us-east-1 Region. Each VPC serves a different business unit and uses its own private VIF for connectivity to the on-premises environment. Users are reporting slowness when they access resources that are hosted on AWS.

A network engineer finds that there are sudden increases in throughput and that the Direct Connect connection becomes saturated at the same time for about an hour each business day. The company wants to know which business unit is causing the sudden increase in throughput. The network engineer must find out this information and implement a solution to resolve the problem.

Which solution will meet these requirements?

A. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observe
B. Create a new 10 Gbps dedicated connectio
C. Shift traffic from the existing dedicated connection to the new dedicated connection.
D. Review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observe
E. Upgrade the bandwidth of the existing dedicated connection to 10 Gbps.
F. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observe
G. Upgrade the existing dedicated connection to a 5 Gbps hosted connection.
H. Review the Amazon CloudWatch metrics for ConnectionBpsIngress and ConnectionPpsEgress to determine which VIF is sending the highest throughput during the period in which slowness is observed.Create a new 10 Gbps dedicated connectio
I. Shift traffic from the existing dedicated connection to the new dedicated connection.

**Answer:** A

**Explanation:**
To meet the requirements of finding out which business unit is causing the sudden increase in throughput and resolving the problem, the network engineer should review the Amazon CloudWatch metrics for VirtualInterfaceBpsEgress and VirtualInterfaceBpsIngress to determine which VIF is sending the highest throughput during the period in which slowness is observed (Option B). After identifying the VIF that is causing the issue, they can upgrade the bandwidth of the existing dedicated connection to 10 Gbps to resolve the problem (Option B).

**NEW QUESTION 64**
A company has deployed an application in a VPC that uses a NAT gateway for outbound traffic to the internet. A network engineer notices a large quantity of suspicious network traffic that is traveling from the VPC over the internet to IP addresses that are included on a deny list. The network engineer must implement a solution to determine which AWS resources are generating the suspicious traffic. The solution must minimize cost and administrative overhead.
Which solution will meet these requirements?

A. Launch an Amazon EC2 instance in the VP
B. Use Traffic Mirroring by specifying the NAT gateway as the source and the EC2 instance as the destinatio
C. Analyze the captured traffic by using open-source tools to identify the AWS resources that are generating the suspicious traffic.
D. Use VPC flow log
E. Launch a security information and event management (SIEM) solution in the VP
F. Configure the SIEM solution to ingest the VPC flow log
G. Run queries on the SIEM solution to identify the AWS resources that are generating the suspicious traffic.
H. Use VPC flow log
I. Publish the flow logs to a log group in Amazon CloudWatch Log
J. Use CloudWatch Logs Insights to query the flow logs to identify the AWS resources that are generating the suspicious traffic.
K. Configure the VPC to stream the network traffic directly to an Amazon Kinesis data strea
L. Send the data from the Kinesis data stream to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Athena to query the data to identify the AWS resources that are generating the suspicious traffic.

**Answer:** C


**NEW QUESTION 66**
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 71**
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 76**
A company is deploying a new application on AWS. The application uses dynamic multicasting. The company has five VPCs that are all attached to a transit gateway Amazon EC2 instances in each VPC need to be able to register dynamically to receive a multicast transmission.
How should a network engineer configure the AWS resources to meet these requirements?

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

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

**Answer:** C


**NEW QUESTION 81**
A media company is implementing a news website for a global audience. The website uses Amazon CloudFront as its content delivery network. The backend runs on Amazon EC2 Windows instances behind an Application Load Balancer (ALB). The instances are part of an Auto Scaling group. The company's customers access the website by using service example com as the CloudFront custom domain name. The CloudFront origin points to an ALB that uses service-alb.example.com as the domain name.
The company's security policy requires the traffic to be encrypted in transit at all times between the users and the backend.
Which combination of changes must the company make to meet this security requirement? (Choose three.)

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

**Answer:** BDE


**NEW QUESTION 83**
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 86**
A company uses a hybrid architecture and has an AWS Direct Connect connection between its on-premises data center and AWS. The company has production applications that run in the on-premises data center. The company also has production applications that run in a VPC. The applications that run in the on-premises data center need to communicate with the applications that run in the VPC. The company is using corp.example.com as the domain name for the on-premises resources and is using an Amazon Route 53 private hosted zone for aws.example.com to host the VPC resources.
The company is using an open-source recursive DNS resolver in a VPC subnet and is using a DNS resolver in the on-premises data center. The company's on-premises DNS resolver has a forwarder that directs requests for the aws.example.com domain name to the DNS resolver in the VPC. The DNS resolver in the

VPC has a forwarder that directs requests for the corp.example.com domain name to the DNS resolver in the on-premises data center. The company has deckled to replace the open-source recursive DNS resolver with Amazon Route 53 Resolver endpoints.

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

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

**Answer:** BCE

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

➢ Create a Route 53 Resolver inbound endpoint and a Route 53 Resolver outbound endpoint. (Option C)

➢ Configure the on-premises DNS resolver to forward aws.example.com domain queries to the IP addresses of the inbound endpoint. (Option B)

➢ Create a Route 53 Resolver rule to forward corp.example.com domain queries to the IP address of the on-premises DNS resolver. (Option E)

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

## NEW QUESTION 87
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 90
A company delivers applications over the internet. An Amazon Route 53 public hosted zone is the
authoritative DNS service for the company and its internet applications, all of which are offered from the same domain name.
A network engineer is working on a new version of one of the applications. All the application's components are hosted in the AWS Cloud. The application has a three-tier design. The front end is delivered through Amazon EC2 instances that are deployed in public subnets with Elastic IP addresses assigned. The backend components are deployed in private subnets from RFC1918.
Components of the application need to be able to access other components of the application within the application's VPC by using the same host names as the host names that are used over the public internet. The network engineer also needs to accommodate future DNS changes, such as the introduction of new host names or the retirement of DNS entries.
Which combination of steps will meet these requirements? (Choose three.)

A. Add a geoproximity routing policy in Route 53.
B. Create a Route 53 private hosted zone for the same domain name Associate the application's VPC with the new private hosted zone.
C. Enable DNS hostnames for the application's VPC.
D. Create entries in the private hosted zone for each name in the public hosted zone by using the corresponding private IP addresses.
E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs when AWS CloudTrail logs a Route 53 API call to the public hosted zon
F. Create an AWS Lambda function as the target of the rul
G. Configure the function to use the event information to update the privatehosted zone.
H. Add the private IP addresses in the existing Route 53 public hosted zone.

**Answer:** BCD

## NEW QUESTION 92
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 95**
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 100**
An ecommerce company is hosting a web application on Amazon EC2 instances to handle continuously changing customer demand. The EC2 instances are part of an Auto Scaling group. The company wants to implement a solution to distribute traffic from customers to the EC2 instances. The company must encrypt all traffic at all stages between the customers and the application servers. No decryption at intermediate points is allowed.
Which solution will meet these requirements?

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

**Answer:** C

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

**NEW QUESTION 105**
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 107**
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 108**
......

# THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual ANS-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the ANS-C01 Product From:

## https://www.2passeasy.com/dumps/ANS-C01/

# Money Back Guarantee

## ANS-C01 Practice Exam Features:

* ANS-C01 Questions and Answers Updated Frequently

* ANS-C01 Practice Questions Verified by Expert Senior Certified Staff

* ANS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* ANS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year