Microsoft Azure AZ-700 Exam Questions

Total Questions: 260+
Demo Questions: 35
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit:
AZ-700 Exam Dumps by Cert Empire

Question: 1

HOTSPOT You have an Azure subscription that contains a dual-stack virtual network named VNet1. VNet1 has the following IP address spaces: • IPv4:192.168.0.0/24 • IPv6: fd0adbftdeca: deed: y48 You plan to deploy an Azure VPN gateway and multiple virtual machines to VNet1. You need to configure the subnet masks for VNet1. The solution must meet the following requirements: • Maximize the number of usable IP addresses. • Support the deployment of the VPN gateway and the virtual machines. Which subnet mask should you use for each address space? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer:

IPv4 subnet mask: /27
IPv6 subnet mask: /64

Explanation:

For the IPv4 address space, an Azure VPN Gateway requires a dedicated subnet, conventionally named GatewaySubnet. Official Azure documentation recommends a GatewaySubnet size of /27 or larger to accommodate various gateway SKUs and potential future scaling. Using a /27 prefix meets this critical requirement while maximizing the number of subnets that can be created within the /24 VNet space, thereby optimizing address space utilization.
For the IPv6 address space, Azure Virtual Network imposes a mandatory requirement that all IPv6 subnets must have a fixed prefix length of /64. This is a non-negotiable constraint for dual-stack network configurations in Azure.

References:

1. Microsoft Learn, Azure VPN Gateway documentation, "Gateway subnet" section. It states, "For the most future-proof configuration, we recommend that you create a gateway subnet of /27 or larger (/27, /26, /25 etc.)." This directly supports the selection of /27 for the IPv4 subnet.
2. Microsoft Learn, "Overview of IPv6 for Azure Virtual Network," Subnetting section. The documentation explicitly states, "The IPv6 subnet size must be /64." This confirms the mandatory prefix length for all IPv6 subnets in an Azure VNet.
3. Microsoft Learn, "Virtual network and subnets," Subnets section. This document explains that a virtual network is segmented into one or more subnets, and each subnet must have a unique address range specified in CIDR format. This provides the foundational context for the subnetting decisions in the question.

Question: 2

HOTSPOT You have an on-premises network. You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
Subnet1	On-premises subnet	Assigned an IP address of 10.1.1.0/24
Subnet2	On-premises subnet	Assigned an IP address of 10.1.2.0/24
Subnet3	Azure virtual subnet	Assigned an IP address of 10.1.3.0/24
Subnet4	Azure virtual subnet	Assigned an IP address of 10.1.1.0/24
VNet1	Azure virtual network	Contains Subnet3 and Subnet4
Server1	Windows Server 2022	On-premises server that is connected to Subnet1 and Subnet2
VM2	Windows Server 2022	Azure virtual machine that is connected to Subnet3 and Subnet4
S2SVPN1	Site-to-Site (S2S) VPN	Connects the on-premises network to VNet1

You need to ensure that on-premises devices can communicate with Azure resources that are connected to Subnet4. What should you do on each resource? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

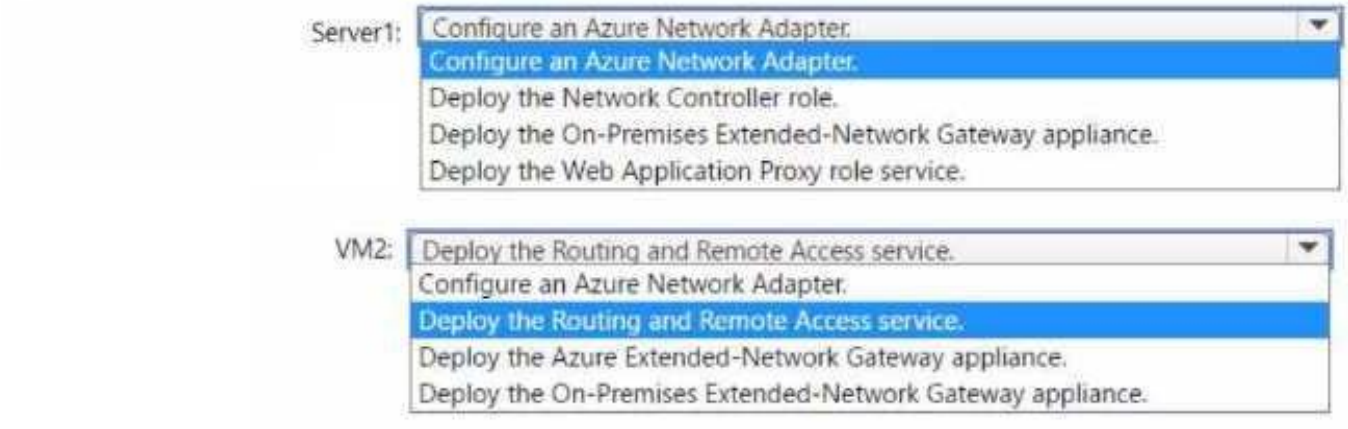

Answer:

Server1: Deploy the On-Premises Extended-Network Gateway appliance.

VM2: Deploy the Routing and Remote Access service.

Explanation:

The core issue is the IP address space overlap between the on-premises Subnet1 (10.1.1.0/24) and the Azure Subnet4 (10.1.1.0/24). A standard Site-to-Site (S2S) VPN cannot route traffic between networks with identical address ranges.

The intended solution is to use Azure Extended Network, a feature designed specifically to stretch an on-premises subnet into Azure, creating a single Layer 2 network segment. This allows virtual

machines in Azure to have IP addresses from the on-premises subnet and communicate as if they were on the local network.

• On Server1 (On-premises): To enable Azure Extended Network, a gateway virtual machine must be deployed in the on-premises environment. The option Deploy the On-Premises Extended-Network Gateway appliance directly corresponds to this requirement. Server1, as an on-premises Windows Server, would act as the host for this gateway.

• On VM2 (Azure): When an Azure VM joins an extended network, its default gateway is reconfigured to point to the on-premises gateway. While this enables communication with the on-premises network, it can isolate the VM from other subnets within its own Azure VNet (like Subnet3). To resolve this, VM2 needs advanced routing capabilities. By deploying the Routing and Remote Access service (RRAS), VM2 can be configured to act as a router. This allows for the creation of static routes to properly direct traffic to both the on-premises network and other Azure resources, overcoming the routing limitations.

References:

Azure Extended Network Documentation (Microsoft Learn): This document explains the feature and its architecture. It explicitly states the need for a gateway VM on-premises to stretch the network.

Reference: Microsoft Corporation. (2023). "Stretch an on-premises subnet into Azure using extended network for Azure". Microsoft Learn. Section: "How it works".

Routing and Remote Access Service (RRAS) Documentation (Microsoft Learn): This documentation details how RRAS turns a Windows Server into a software router, capable of managing complex traffic flows with static routes, which is what is needed on VM2 to handle the new routing complexity.

Reference: Microsoft Corporation. (2021). "Routing and Remote Access Service (RRAS)". Microsoft Learn. Section: "Software-defined networking (SDN) router".

Question: 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure subscription that contains an Azure Virtual WAN named VWAN1. VWAN1 contains a hub named Hub1. Hub1 has a security status of Unsecured. You need to ensure that the security status of Hub1 is marked as Secured. Solution: You implement Azure Web Application Firewall (WAF). Does this meet the requirement?

A. Yes

B. No

Answer:

B

Explanation:

The security status of an Azure Virtual WAN hub changes from "Unsecured" to "Secured" only when a network security service, such as Azure Firewall or a supported third-party Network Virtual Appliance (NVA), is deployed within the hub via Azure Firewall Manager. This action converts the standard hub into a "Secured Virtual Hub."
Azure Web Application Firewall (WAF) is a Layer 7 security service designed to protect web applications from common vulnerabilities. It is deployed with services like Azure Application Gateway or Azure Front Door, not directly within a Virtual WAN hub to change its fundamental security status. Therefore, implementing WAF does not meet the requirement.

Why Incorrect Options are Wrong:

A. Yes: This is incorrect. Azure WAF protects web applications and does not fulfill the specific requirement of deploying a network firewall within the Virtual WAN hub to change its status to "Secured."

References:

1. Microsoft Learn Azure Firewall Manager What is a secured virtual hub?
Section: "What is a secured virtual hub?"
Content: "A secured virtual hub is an Azure Virtual WAN Hub with associated security and routing policies configured by Azure Firewall Manager. Use secured virtual hubs to easily create hub-and-spoke and transitive architectures with native security services for traffic governance and protection." This document explicitly states that a hub becomes "secured" by integrating security services like Azure Firewall via Firewall Manager.

2. Microsoft Learn Azure Web Application Firewall What is Azure Web Application Firewall?
Section: "Overview"
Content: "Azure Web Application Firewall (WAF) provides centralized protection of your web applications from common exploits and vulnerabilities. WAF can be deployed with Azure Application Gateway, Azure Front Door, and Azure Content Delivery Network (CDN) service from Microsoft." This reference clarifies that WAF's purpose and deployment model are distinct from securing a Virtual WAN hub's infrastructure.
3. Microsoft Learn Azure Virtual WAN Create a secured virtual hub using Azure portal
Section: "Prerequisites" and "Create a secured virtual hub"
Content: The tutorial steps demonstrate that creating a secured virtual hub involves selecting the "Include gateway" option and then configuring Azure Firewall within the hub's settings. This confirms that Azure Firewall is the required component, not WAF.

Question: 4

HOTSPOT You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
VNet1	Virtual network	In the West Europe Azure region
VNet2	Virtual network	In the East US Azure region
VM1	Virtual machine	On VNet1
VM2	Virtual machine	On VNet1
VM3	Virtual machine	On VNet2
VM4	Virtual machine	On VNet2

You plan to deploy an app named App1 to meet the following requirements. • External users must be able to access App1 from the internet. • App1 will be load balanced across all the virtual machines. • App1 will be hosted on VM1, VM2. VM3. and VM4. • App1 must be available if an Azure region fails. • Costs must be minimized. You need to implement a global load balancer solution for App. What should you configure? To answer, select the appropriate options in the answer area NOTE: Bach correct answer is worth one point.

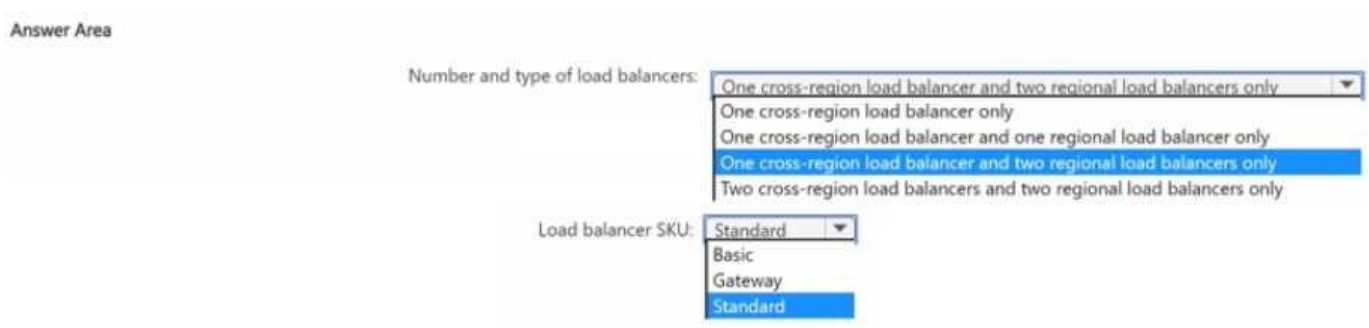

Answer:

Number and type of load balancers: One cross-region load balancer and two regional load balancers only Load balancer SKU: Standard

Explanation:

The scenario requires a load-balanced solution for an application (App1) deployed across virtual machines in two different Azure regions (West Europe and East US). A key requirement is that the application must remain available even if an entire Azure region fails.

To meet this requirement for high availability across regions, a global load balancing solution is necessary. The Azure Cross-region Load Balancer is specifically designed for this purpose.

The architecture for a cross-region load balancer consists of:

• A single cross-region load balancer with a static, global anycast public IP address to act as the primary entry point for traffic.

• A backend pool for the cross-region load balancer that contains the front-end configurations of regional load balancers.

In this case, one regional load balancer is needed in the West Europe region to manage traffic for VM1 and VM2, and a second regional load balancer is needed in the East US region for VM3 and VM4. This results in a total of one cross-region load balancer and two regional load balancers.

Regarding the SKU, the cross-region load balancer functionality is only available with the Standard SKU. Furthermore, any regional load balancers added to the backend pool of a cross-region load balancer must also be of the Standard SKU. The Basic SKU does not support cross-region deployments and lacks the necessary availability features.

References:

Microsoft Azure Documentation, "What is Cross-region Load Balancer?". Microsoft Learn, Retrieved September 19, 2025.

Section: Why use Cross-region load balancer?: "Cross-region load balancer provides a static global anycast public IP address for your multi-region application...If a region fails, you can direct traffic to the next closest healthy region."

Section: Architecture: "The backend pool of a cross-region load balancer contains one or more regional load balancers."

Section: Cross-region load balancer and regional load balancer: "A cross-region load balancer is a Standard SKU public load balancer...The regional load balancers you add to the backend of the cross-region load balancer must be a Standard SKU load balancer."

Microsoft Azure Documentation, "Azure Load Balancer SKUs". Microsoft Learn, Retrieved September 19, 2025.

Table: SKU comparison: This table explicitly states that the Basic SKU has "No SLA" and is not recommended for production workloads, while the Standard SKU has a 99.99% SLA and supports features like Availability Zones, which are critical for high-availability scenarios. It also notes that cross-region load balancing is a feature of the Standard SKU.

Question: 5

You have an Azure Private Link service named PL1 that uses an Azure load balancer named LB1. You need to ensure that PL1 can support a higher volume of outbound traffic. What should you do?

A. Redeploy LB1 with a different SKU.

B. Increase the number of NAT IP addresses assigned to PL1.

C. Deploy an Azure Application Gateway v2 instance to the source NAT subnet.

D. Increase the number of frontend IP configurations for LB1.

Answer:

B

Explanation:

The Azure Private Link service uses Source Network Address Translation (SNAT) for traffic originating from the service provider's virtual network and destined for the private endpoint in the consumer's network. The source IP of these packets is translated to a NAT IP address allocated from the provider's VNet. To support a higher volume of outbound traffic and avoid SNAT port exhaustion, you must scale the number of available SNAT ports. This is achieved by increasing the number of NAT IP addresses configured for the Private Link service. Each additional NAT IP address provides more available ports for translation.

Why Incorrect Options are Wrong:

A. The load balancer SKU (Standard is required for Private Link) does not control the dedicated outbound NAT capacity of the Private Link service itself.
C. An Azure Application Gateway is a Layer 7 web traffic load balancer and is not used to scale the outbound NAT of a Layer 4 Private Link service.
D. Frontend IP configurations on the load balancer are for receiving inbound traffic, not for scaling the Private Link service's outbound NAT capabilities.

References:

1. Microsoft Learn What is Azure Private Link service?: Under the section "Source Network Address Translation (SNAT)", it states, "To scale, add more NAT IP addresses to your Private Link service. Each new NAT IP address adds 64,000 more available SNAT ports."
2. Microsoft Learn Troubleshoot Azure Private Link service connectivity problems: In the section "Private Link service has SNAT port exhaustion", the recommended solution is: "The Private Link service has up to 8 NAT IPs that can be used for SNAT. Each NAT IP can assign 64k ports. Add more NAT IPs to the Private Link service to avoid SNAT port exhaustion."

Question: 6

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains the following subnets. • AzureFirewallSubnet • GatewaySubnet • Subnet 1 • Subnet2 • Subnet3 Subnet2 has a delegation to the Microsoft.Web/serverfarms service. The subscription contains the resources shown in the following table.

Name	Type	Connected to
AZVNGW1	Azure VPN Gateway	GatewaySubnet
AZFW1	Azure Firewall Premium	AzureFirewallSubnet
VMSS1	Virtual machine scale set	Subnet1

You need to implement an Azure application gateway named AG1 that will be integrated with an Azure Web Application Firewall (WAF). AG1 will be used to publish VMSS1. To which subnet should you connect AG1?

A. Subrwt2

B. Subnet 1

C. Subnet3

D. AzureFjrewall Subnet

E. GatewaySubnet

Answer:

C

Explanation:

Azure Application Gateway requires a dedicated subnet within the virtual network. This subnet cannot be shared with any other resources.
GatewaySubnet and AzureFirewallSubnet are reserved names for their specific gateway and firewall services, respectively, and cannot be used for an Application Gateway.
Subnet2 is delegated to the Microsoft.Web/serverfarms service, which means it can only host resources of that type and is therefore unavailable.
Subnet1 already contains the VMSS1 resource, so it cannot be used as the dedicated subnet for the Application Gateway.
This leaves Subnet3 as the only available and valid option, as it is not reserved, not delegated, and does not contain other resources.

Why Incorrect Options are Wrong:

A. Subnet2: This subnet is delegated to the Microsoft.Web/serverfarms service and cannot be used for deploying an Application Gateway.
B. Subnet1: This subnet already contains the VMSS1 resource. An Application Gateway requires

its own dedicated subnet.
D. AzureFirewallSubnet: This is a reserved subnet name and can only be used for deploying an Azure Firewall.
E. GatewaySubnet: This is a reserved subnet name and can only be used for deploying virtual network gateways (VPN or ExpressRoute).

References:

1. Microsoft Learn Azure Application Gateway infrastructure configuration Subnet: "The application gateway needs a dedicated subnet. The subnet can contain only application gateways. No other resources are allowed in the subnet."
2. Microsoft Learn What is subnet delegation? Constraints on delegation: "You can't deploy resources from different services into a delegated subnet." This rule disqualifies Subnet2.
3. Microsoft Learn GatewaySubnet: "When you create a virtual network gateway, you must specify the gateway subnet that you want to deploy it to... Don't deploy any other resources (for example, VMs) to the gateway subnet." This rule disqualifies GatewaySubnet.
4. Microsoft Learn Azure Firewall FAQ Why does Azure Firewall need a dedicated subnet?: "Azure Firewall must be deployed in a dedicated subnet. It can't be deployed in a subnet with other resources. The name of this subnet must be AzureFirewallSubnet." This rule disqualifies AzureFirewallSubnet.

Question: 7

DRAG DROP You have an Azure subscription that contains the resources shown in the following table. You discover that users connect directly to App1. You need to meet The following requirements: • Administrators must only access App1 by using a private endpoint. • All user connections to App1 must be routed through FD1. • The downtime of connections to App1 must be minimized. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

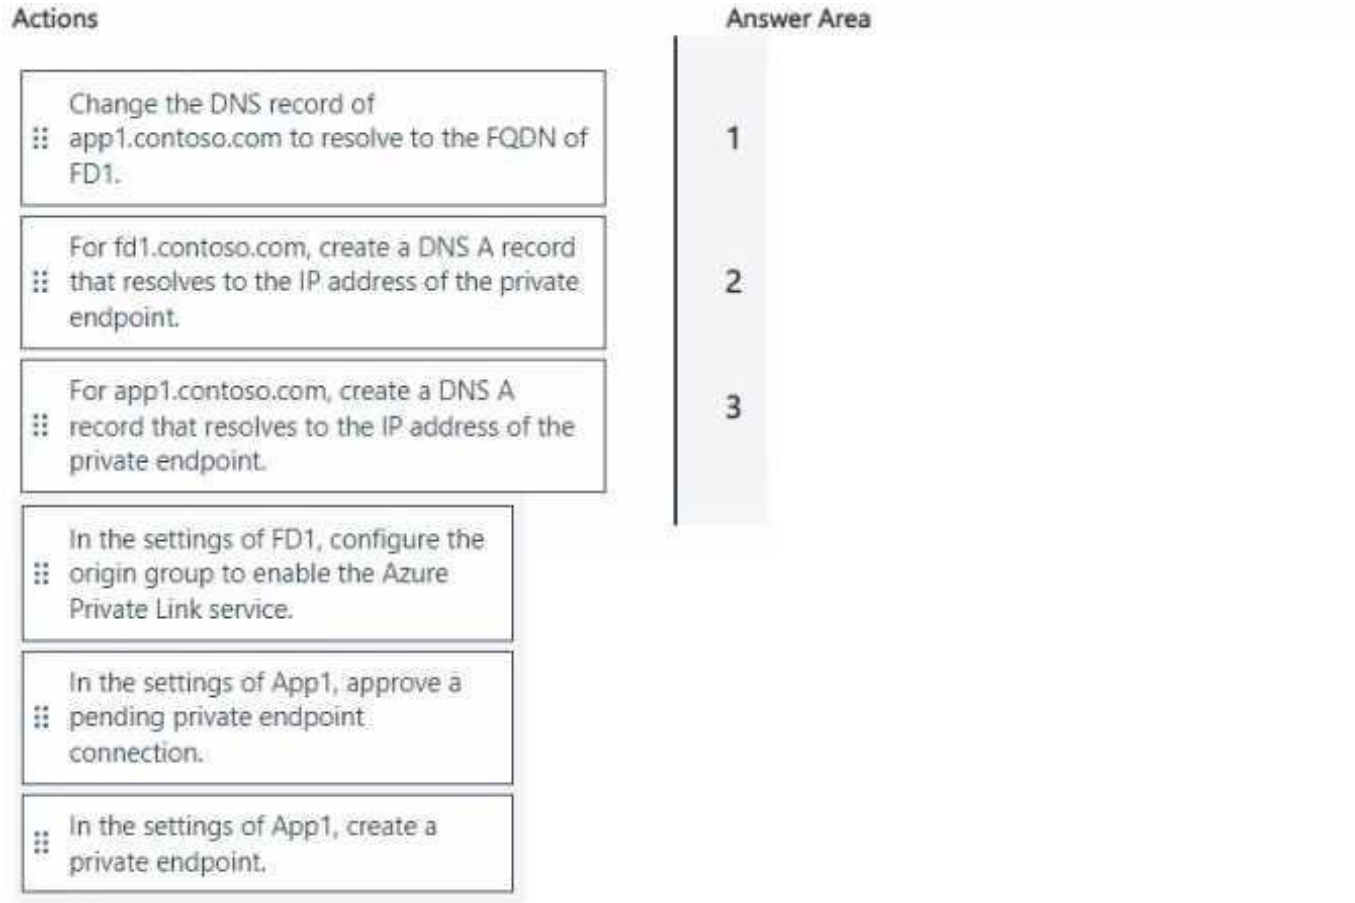

Answer:

1: In the settings of FD1, configure the origin group to enable the Azure Private Link service.

2: In the settings of App1, approve a pending private endpoint connection.

3: Change the DNS record of https://www.google.com/url?sa=E&source=gmail&q=app1.contoso.com to resolve to the FQDN of FD1.

Explanation:

This sequence correctly establishes a new, private traffic path from Azure Front Door (FD1) to the App Service (App1) before redirecting live user traffic, which is crucial for minimizing downtime.

• Configure FD1 to use Private Link: The first step is to configure the origin group in the Azure Front Door Premium instance to connect to the App Service origin using the Azure Private Link service. This action initiates a request to create a private endpoint connection on the App Service.

• Approve the connection on App1: The private endpoint connection created by Front Door will be in a "Pending" state. You must navigate to the App Service's networking settings and approve this pending connection. This completes the setup of the secure, private path between Front Door and the App Service. At this stage, the new path is operational, but user traffic is still going directly to App1.

• Change the DNS record: The final step is to perform the traffic cutover. By changing the public DNS CNAME record for app1.contoso.com to point to the FQDN of the Front Door instance (FD1), you redirect all user traffic to flow through Front Door. Since the private backend connection is already established, there is no interruption in service.

This order ensures the new infrastructure is fully functional before it receives live traffic, fulfilling the requirement to minimize downtime.

References:

Microsoft Learn, Azure Front Door Documentation: The official documentation on securing origin with Private Link in Azure Front Door Premium outlines this exact procedure. It specifies that you first enable Private Link on the Front Door origin, which creates a pending private endpoint connection that must then be approved on the origin resource.
Reference: Secure your Origin with Private Link in Azure Front Door Premium. See the sections "Enable Private Link to an App service" and "Approve the private endpoint connection from the web app." This guide confirms the sequence of enabling the service in Front Door first, followed by approving the connection on the App Service.
Microsoft Learn, Tutorial: Connect to a web app using an Azure Private Endpoint: This tutorial provides general information on private endpoints for web apps. It explains that "When you create a private endpoint... a connection approval workflow is required." This supports the necessity of the approval step after the connection is initiated.
Reference: Connect to a web app using an Azure Private Endpoint. Refer to the "Create a private endpoint" section, which details the approval process.

Question: 8

You have an instance of Azure Web Application Firewall (WAF) on Azure Front Door. You plan to create a WAF rule that will block high rates of requests from a single IP address. You need to query Log Analytics to identify the optimal threshold for the rule. Which table should you query in Log Analytics?

A. AZFWThreatInte1

B. AzureDiagnostics

C. SecurityDetection

D. AGWFirewallLogs

Answer:

B

Explanation:

Azure Web Application Firewall (WAF) on Azure Front Door integrates with Azure Monitor to store its logs in a Log Analytics workspace. When using the classic Azure Diagnostics settings, these logs are sent to the AzureDiagnostics table. This table contains detailed information for each request evaluated by the WAF, including the clientIps field, which records the source IP address. To determine an optimal threshold for a rate-limiting rule, you can execute a Kusto Query Language (KQL) query against this table to count the number of requests per clientIps over a specific time interval.

Why Incorrect Options are Wrong:

A. AZFWThreatInte1: This table stores threat intelligence logs specifically for the Azure Firewall service, not for WAF on Azure Front Door.
C. SecurityDetection: This table is used by Microsoft Defender for Cloud to store security alerts and detections, not the raw request logs needed for rate analysis.
D. AGWFirewallLogs: This is a resource-specific table for logs from WAF on Azure Application Gateway (AGW), not Azure Front Door.

References:

1. Microsoft Learn, Azure Web Application Firewall documentation. "Azure Web Application Firewall monitoring and logging". This document explicitly states, "Log Analytics logs are stored in the AzureDiagnostics table." It also provides a sample query for WAF logs that queries the AzureDiagnostics table.
Reference: Section "Logs and metrics", Paragraph 2.
2. Microsoft Learn, Azure Front Door documentation. "Data reference for monitoring Azure Front Door". This page details the schema for Azure Front Door logs sent to Log Analytics. For the

FrontdoorWebApplicationFirewallLog category, it lists the columns available in the AzureDiagnostics table, including clientIps, which is essential for the task described.
Reference: Section "AzureDiagnostics table", Subsection "FrontdoorWebApplicationFirewallLog".
3. Microsoft Learn, Azure Firewall documentation. "Azure Firewall logs and metrics". This document describes the tables used for Azure Firewall logging, confirming that AZFWThreatIntel is for threat intelligence logs from Azure Firewall.
Reference: Section "Log Analytics tables", Row "AZFWThreatIntel".
4. Microsoft Learn, Azure Application Gateway documentation. "Resource logs for Azure Application Gateway". This source details the logging options for Application Gateway, specifying that AGWFirewallLogs is the resource-specific table for its WAF logs.
Reference: Section "Resource-specific tables", Table "Resource-specific".

Question: 9

You have the Azure subscriptions shown in the following table.

Name	Microsoft Entra ID tenant	Contains resources in Azure region	Virtual network
Sub1	contoso.com	East US, West US	VNet1, VNet2
Sub2	contoso.com	Europe North, Europe West	VNet3, VNet4
Sub3	fabrikam.com	Europe North, West US	VNet5, VNet6

Each virtual network contains 20 internet-accessible resources that are assigned public IP addresses. You need to implement Azure DDoS Network Protection to protect the resources. The solution must minimize costs. What is the minimum number of DDoS Network Protection plans you should deploy?

A. 1

B. 2

C. 3

D. 6

Answer:

B

Explanation:

An Azure DDoS Network Protection (formerly DDoS Standard) plan is billed once per plan and can be linked to as many as 10 virtual networks that reside in subscriptions belonging to the same Azure AD tenant.

Virtual networks that belong to a different tenant cannot be linked to that plan; a separate plan is required for each tenant.

The table shows the six virtual networks are spread across two different Azure AD tenants.

Therefore, to cover all 40 public IP resources while minimising cost, you need exactly one plan per tenant, i.e., two plans in total.

Why Incorrect Options are Wrong:

A. 1 - A single plan cannot be linked to virtual networks in two different Azure AD tenants.

C. 3 - No technical or billing requirement forces a third plan; tenants, not regions or subscriptions, dictate the minimum.

D. 6 - One plan per virtual network is unnecessary; multiple VNets in the same tenant can share a single plan.

References:

1. Microsoft Azure documentation, "DDoS protection plans - associate or dissociate a virtual network," learn.microsoft.com/en-us/azure/ddos-protection/manage-ddos-protection?tabs=azure-portal#associate-a-virtual-network (see "All virtual networks must belong to subscriptions that are associated with the same Azure Active Directory tenant").
2. Microsoft Azure documentation, "Azure DDoS Protection Standard overview," learn.microsoft.com/en-us/azure/ddos-protection/ddos-protection-overview#pricing (section "Billing and limits" - fixed monthly charge per plan; plan can protect multiple VNets).
3. Microsoft Azure documentation, "Best practices for Azure DDoS Protection," learn.microsoft.com/en-us/azure/ddos-protection/ddos-best-practices (paragraph "Reuse existing plans across virtual networks in the same tenant").

Question: 10

HOTSPOT You plan to implement an Azure Virtual WAN named VWAN1 that will contain a hub named Hub1. VWAN1 will include the virtual networks shown in the following table.

Name	IP address space	Description
VNet1	10.1.0.0/24	Connected directly to Hub1 by using a connection named Conn1
VNet2	10.2.0.0/24	Connected directly to Hub1 by using a connection named Conn2 and hosting a Network Virtual Appliance (NVA) named NVA2 that has an IP address of 10.2.0.5
VNet3	10.2.3.0/24	Connected to VNet2 by using a virtual network peering named Peering1

You need to ensure that hosts connected to VNet1 can communicate with hosts connected to VNet3. How should you configure the routing tables for VWAN1? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

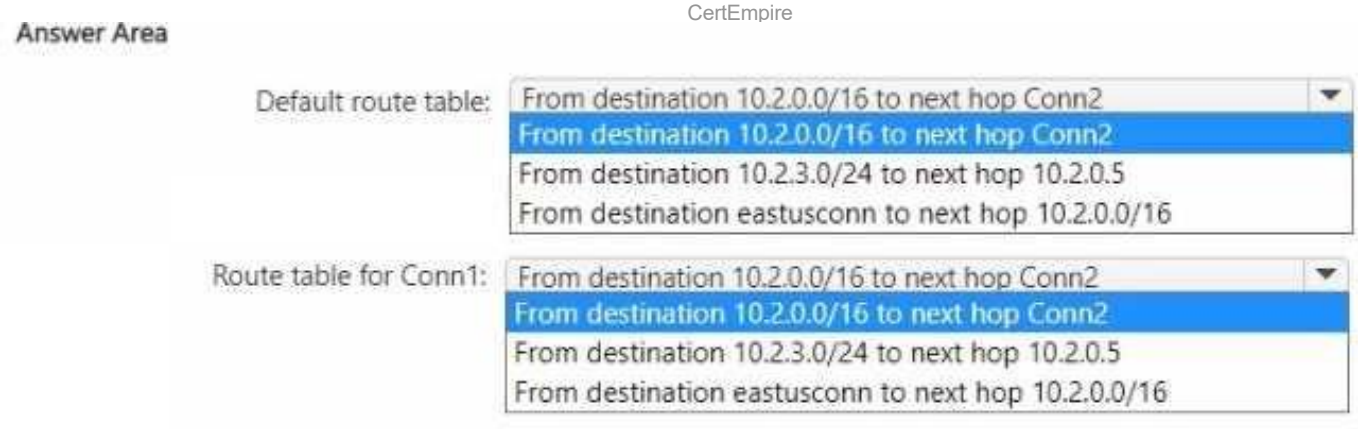

Answer:

Default route table: From destination 10.2.3.0/24 to next hop 10.2.0.5

Route table for Conn1: From destination 10.2.3.0/24 to next hop 10.2.0.5

Explanation:

To enable communication from VNet1 to VNet3, traffic must be routed through the Network Virtual Appliance (NVA) located in VNet2.

• Traffic originating from VNet1 arrives at the virtual hub (Hub1) via its connection, Conn1.

• The hub must have a route to direct traffic destined for VNet3's address space (10.2.3.0/24) to the NVA's IP address (10.2.0.5).

• This is achieved by adding a static route to the hub's routing table. The Default route table is used by all connected VNets by default. Therefore, the static route (10.2.3.0/24 - 10.2.0.5) must be added to it.

• The "Route table for Conn1" represents the routing context for that specific connection. It must also contain this route to ensure traffic from VNet1 is forwarded correctly. Since the other options are invalid or would bypass the NVA, the same static route is the only correct choice.

References:

Microsoft Docs, Azure Virtual WAN: "How to configure virtual hub routing". This document explains how to add static routes to a virtual hub's route table. In the section "To add/update a static route," it specifies that the "Next hop IP address" should be the IP address of the NVA.
Microsoft Docs, Azure Virtual WAN Scenarios: "Route traffic through an NVA". This official documentation describes the exact scenario presented. It confirms that to route traffic from a spoke VNet to another network via an NVA, a static route must be configured on the virtual hub's route table. The route's destination prefix is the target network (10.2.3.0/24), and the next hop is the IP address of the NVA (10.2.0.5) in the connected VNet.

Question: 11

HOTSPOT You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
VNet1	Virtual network	Has an IP address space of 192.168.0.0/23
VNet2	Virtual network	Has an IP address space of 192.168.2.0/23
VNet3	Virtual network	Has an IP address space of 10.0.0.0/20
Peering12	Virtual network peering	Peered between VNet1 and VNet2
Peering21	Virtual network peering	Peered between VNet2 and VNet1

Each virtual network contains 20 virtual machines and a subnet that has an IP address space of /24. You need to ensure that you can access the virtual machines from the internet by using Azure Bastion. What is the minimum number of bastion subnets you should deploy, and what is the smallest supported IP address space for each bastion subnet? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Bastion subnets: 2

IP address space: /26

Explanation:

Azure Bastion is deployed within a virtual network (VNet) and provides secure access to virtual machines in that VNet, as well as in directly peered VNets.

• Number of Bastion Subnets: Since VNet1 and VNet2 are peered, a single Azure Bastion host deployed in either VNet1 or VNet2 can be used to access VMs in both networks. VNet3 is not peered with the other two, so it requires its own separate Azure Bastion host to enable access to its VMs. Therefore, a minimum of two Bastion deployments, and consequently two Bastion subnets, are needed.

• IP Address Space: Azure documentation specifies that for a Bastion deployment, a dedicated subnet named AzureBastionSubnet is required. The smallest supported address space for this subnet is a /26 prefix. While larger subnets (e.g., /25, /24) are valid, the question asks for the smallest supported size.

References:

Microsoft Learn, Azure Bastion Documentation, "About VNet peering and Azure Bastion": "If you have VNet peering configured, you don't have to deploy Azure Bastion in each peered VNet. This means if you have an Azure Bastion host configured in one virtual network (VNet), you can use it to connect to VMs deployed in a peered VNet without deploying an additional Bastion host." Microsoft Learn, Azure Bastion Documentation, "Tutorial: Create an Azure Bastion host": In the "Prerequisites" section, it states, "The virtual network must have a dedicated subnet named AzureBastionSubnet. This subnet must be created with a /26 prefix or larger (/25, /24 etc.)." Microsoft Learn, Azure Bastion Documentation, "Azure Bastion FAQ": "What are the required NSG rules for the Azure Bastion subnet? ... Does Azure Bastion require an AzureBastionSubnet of size /27 or larger? Azure Bastion requires a dedicated subnet, AzureBastionSubnet, of size /26 or larger."

Question: 12

HOTSPOT You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Location	Description
SQLMI1	Azure SQL Managed Instance	US East	Managed instance connected to VNet1
contoso.com	Microsoft Entra Domain Services	US East	Domain connected to VNet2
VNet1	Virtual network	US East	None
VNet2	Virtual network	US East	None
storage1	Storage account	US East	None

You need to ensure that network traffic is routed over the Azure backbone network for the following scenarios: • Traffic from SQIMI1 to storage1 • Traffic from domain joined servers on VNet2 to storage1 The solution must minimize costs. What should you configure for each scenario? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

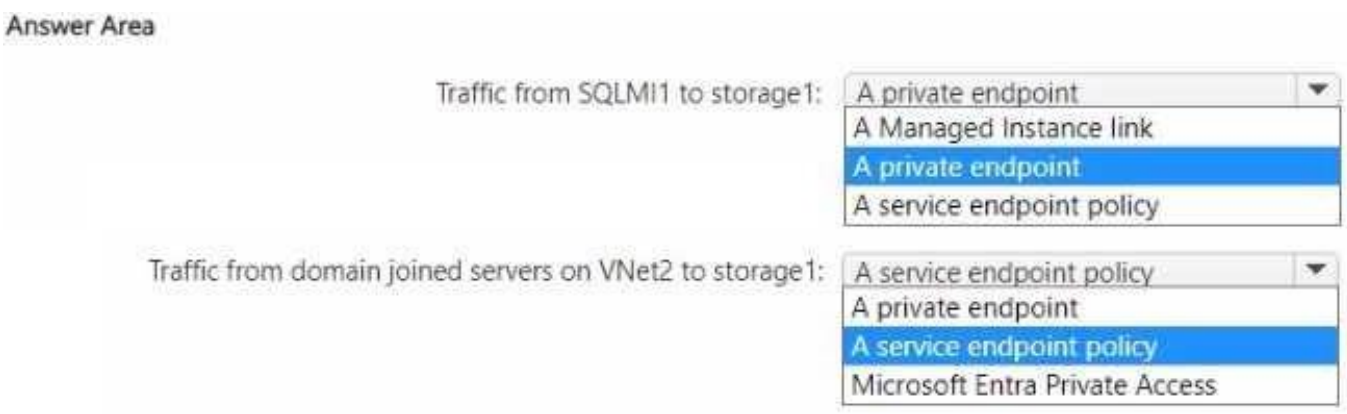

Answer:

Traffic from SQLMI1 to storage1: A private endpoint

Traffic from domain joined servers on VNet2 to storage1: A service endpoint policy

Explanation:

For traffic from SQLMI1 to storage1, a private endpoint is the correct solution. Azure SQL Managed Instance is deployed into a dedicated subnet with specific network configurations. Crucially, subnets hosting a SQL Managed Instance do not support service endpoints for connectivity from the instance to Azure Storage. Therefore, to securely route traffic to storage1 over the Azure backbone, a private endpoint must be created for the storage account within

VNet1. This makes the storage account accessible via a private IP address in the virtual network, satisfying the requirement.

For traffic from domain-joined servers on VNet2 to storage1, a service endpoint policy is the most precise and cost-effective solution. A virtual network service endpoint for Azure Storage routes traffic from the VNet directly to the storage service on the Azure backbone, and this feature is free, meeting the cost-minimization requirement. A service endpoint policy is configured on top of the service endpoint to explicitly allow access only to the storage1 account, preventing data exfiltration to other storage accounts. This is more specific and secure than using a service endpoint alone.

References:

Microsoft Learn. Connectivity architecture for Azure SQL Managed Instance. In the "High-level connectivity architecture" section, it is stated: "Connections from SQL Managed Instance to Azure Storage or Azure Key Vault don't work if they use service endpoints." This confirms why a private endpoint is the only functional choice for the SQLMI1 scenario.
Microsoft Learn. Virtual Network service endpoints. Under the "Pricing and limits" section, it confirms the cost-effectiveness: "There's no additional charge for using service endpoints." This supports its selection for the VNet2 scenario where cost is a factor.
Microsoft Learn. Virtual network service endpoint policies for Azure Storage. This document explains that policies "allow you to filter egress virtual network traffic to specific Azure Storage accounts over a service endpoint". This aligns with configuring access specifically for storage1, making it a more precise answer than a service endpoint alone.
Azure Pricing. Azure Private Link pricing. This page details the costs associated with private endpoints, including an hourly resource charge and a per-GB data processing charge, confirming they are not the most cost-effective option when a free alternative (service endpoint) is available and supported.

Question: 13

You have an Azure subscription. You plan to implement Azure Virtual WAN as shown in the following exhibit. What is the minimum number of route tables that you should create?

A. 1

B. 2

C. 4

D. 6

Answer:

B

Explanation:

When an Azure Virtual WAN hub is created, it automatically includes two built-in route tables: 'Default' and 'None'. The 'Default' route table is configured to learn routes from all connected virtual networks, sites, and point-to-site users, enabling any-to-any connectivity by default. The 'None' route table can be used to isolate specific connections by preventing route propagation. For the topology shown in the exhibit, which depicts standard connectivity without any specific isolation requirements, these two automatically created route tables are sufficient. Therefore, the minimum number of route tables that must exist is two.

Why Incorrect Options are Wrong:

A: A virtual hub is always provisioned with two route tables ('Default' and 'None'), so a single route table is not a valid minimum.
C: Four route tables would only be necessary for advanced routing scenarios requiring multiple, distinct isolation domains, which is not the minimum.
D: Six route tables would be for highly complex, segmented routing scenarios and is far more than the minimum required.

References:

1. Microsoft Learn Azure Virtual WAN documentation. In the article "About virtual hub routing," the "Route tables" section explicitly states: "When a virtual hub is created, a Default route table and a None route table is created in the virtual hub." This confirms that two tables are the baseline. Source: Microsoft Learn, "About virtual hub routing", Section: "Route tables".
2. Microsoft Learn Azure Virtual WAN documentation. The same document further explains the purpose of the 'Default' route table, which supports the scenario in the exhibit: "The Default route table has a built-in mechanism that automatically learns routes... By default all connections...propagate routes to the Default route table." This shows that no additional tables are needed for the basic connectivity shown.

Source: Microsoft Learn, "About virtual hub routing", Section: "Default route table".

Question: 14

DRAG DROP You have a computer named CLIENT! that runs Windows 11 and has the Azure VPN Client installed. You have an Azure virtual network gateway named VPNGW1. You need to ensure that you can connect CLIENT1 to VPNGW1. The solution must support Microsoft Entra authentication. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

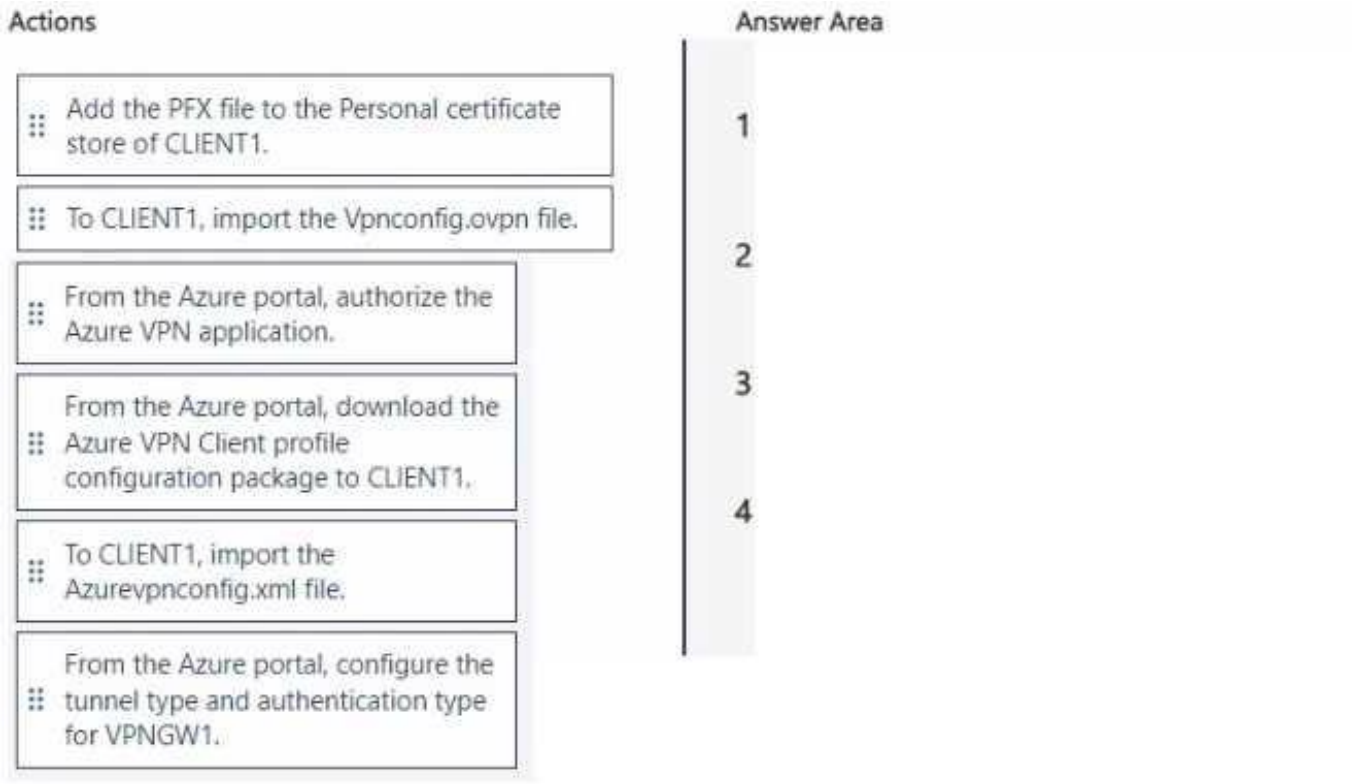

Answer:

From the Azure portal, configure the tunnel type and authentication type for VPNGW1.

From the Azure portal, authorize the Azure VPN application.

From the Azure portal, download the Azure VPN Client profile configuration package to CLIENT1.

To CLIENT1, import the Azurevpnconfig.xml file.

Explanation:

The correct sequence involves first configuring the server-side components in Azure and then configuring the client machine.

• Configure Gateway: The initial step is to configure the Point-to-Site (P2S) settings on the virtual network gateway (VPNGW1) in the Azure portal. You must specify the tunnel type (OpenVPN) and set the authentication type to Microsoft Entra ID.

• Authorize Application: For Microsoft Entra authentication to work, an administrator must grant tenant-wide admin consent to the "Azure VPN" enterprise application. This authorization allows the VPN application to sign in and read user profiles.

• Download Profile: After the gateway is configured and the application is authorized, you generate and download the Azure VPN Client profile configuration package. This package contains the necessary settings, including the azurevpnconfig.xml file.

• Import Configuration: The final step is to configure the client. You must import the azurevpnconfig.xml file from the downloaded package into the Azure VPN Client application on CLIENT1. This action creates the VPN connection profile on the client.

References:

Microsoft Learn: The article "Configure a Point-to-Site VPN connection using Microsoft Entra authentication" outlines these steps precisely.
Step 1 (Configure Gateway): Detailed under the section "2. Configure P2S on the gateway." This section specifies setting the Tunnel type to OpenVPN (SSL) and the Authentication type to Microsoft Entra ID.
Step 2 (Authorize Application): Covered in the section "Enable Microsoft Entra authentication on the virtual network gateway" under "1. Enable authentication." It requires providing admin consent to the "Azure VPN" application.
Step 3 (Download Profile): Explained in the section "3. Download the Azure VPN Client profile configuration package."
Step 4 (Import Configuration): Described in the section "4. Configure the Azure VPN Client," which explicitly states to import the downloaded azurevpnconfig.xml file.

Question: 15

You have an on-premises datacenter named Site1 that contains a firewall named FW1. FW1 connects to the internet. You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
VNet1	Virtual network	*None*
VWAN1	Azure Virtual WAN	Standard Virtual WAN connected to Hub1
Hub1	Azure Virtual WAN hub	Contains a Site-to-Site (S2S) VPN gateway

You plan to connect Site1 to Hub1 by using a site-to-site connection. You need to configure the site-to-site connection to FW1. What should you create in VWAN1?

A. a VPN site

B. a virtual network connection

C. a network virtual appliance (NVA)

D. a User VPN configuration

Answer:

A

Explanation:

To establish a site-to-site (S2S) VPN connection from an on-premises location to an Azure Virtual WAN hub, you must first create a resource that logically represents the on-premises VPN device and its location. In Azure Virtual WAN, this representative object is called a VPN site. The VPN site contains the configuration details for the on-premises device, such as its public IP address, BGP settings, and link information. Once the VPN site is created, it is then associated with the virtual hub to establish the S2S VPN tunnels.

Why Incorrect Options are Wrong:

B. a virtual network connection: This is used to connect an Azure Virtual Network (VNet) to a Virtual WAN hub, not for connecting an on-premises site.
C. a network virtual appliance (NVA): An NVA is a virtual machine that provides network functions within Azure. It does not represent an on-premises location for a VPN connection.
D. a User VPN configuration: This is used to configure point-to-site (P2S) VPNs for individual client devices, not for a site-to-site connection between datacenters.

References:

1. Microsoft Learn. (2023). Tutorial: Create a Site-to-Site connection using Azure Virtual WAN. In the "Prerequisites" section, it states, "You have an on-premises VPN device... This is referred to as a VPN site." Step 2 of the tutorial is explicitly "Create a site," which involves creating the VPN site resource to represent the on-premises location.
2. Microsoft Learn. (2023). Azure Virtual WAN FAQ. Under the "Virtual WAN concepts" section, the question "What is a Virtual WAN site?" is answered with: "A Virtual WAN site is the same as a VPN site. It represents your on-premises location."
3. Microsoft Learn. (2023). Connect a virtual network to a Virtual WAN hub. This document states, "This article helps you connect your virtual network to your virtual hub using a virtual network connection." This confirms that virtual network connections are for VNet-to-hub connectivity, not on-premises sites.
4. Microsoft Learn. (2023). Configure a User VPN (point-to-site) connection for Virtual WAN. The introduction clearly defines this feature: "This article shows you how to use Virtual WAN to connect to your resources in Azure over an IKEv2 or OpenVPN IPsec/VPN connection. This type of connection requires a VPN client to be configured on the client computer." This distinguishes it from a site-to-site connection.

Question: 16

You have an on-premises network named Site1. You have an Azure subscription that contains a storage account named storage1 and a virtual network named VNet1. VNet1 contains a subnet named Subnet1. A private endpoint for storage1 is connected to Subnet1 Site1 is connected to VNet1 by using a Site-to-Site (S2S) VPN. You need to control access to storage1 from Site1 by using network security groups (NSGs). What should you do first?

A. Associate a route table with Subnet1.

B. Associate a NAT gateway with Subnet1.

C. Configure a network policy for private endpoints on Subnet1.

D. Create a subnet delegation on Subnet1.

Answer:

C

Explanation:

To control network traffic to a private endpoint using a Network Security Group (NSG), the network policy for private endpoints must be enabled on the subnet where the private endpoint is located. By default, this policy is disabled, which means NSGs and User-Defined Routes (UDRs) do not apply to traffic destined for private endpoints within that subnet. Enabling this policy is the necessary first step to allow the NSG associated with Subnet1 to inspect and filter traffic, including the traffic originating from the on-premises network Site1.

Why Incorrect Options are Wrong:

A. Associating a route table controls traffic routing but does not enable NSG filtering for private endpoints; the network policy must be enabled for NSGs to take effect.
B. A NAT gateway is used for outbound internet connectivity from a subnet and is not relevant for controlling inbound traffic to a private endpoint from an on-premises network.
D. Subnet delegation is used to dedicate a subnet for a specific PaaS service to inject its resources, which is unrelated to applying NSG rules to a private endpoint.

References:

1. Microsoft Learn Azure Private Link Documentation. "Manage network policies for private endpoints." This document explicitly states, "To apply a network security group to a private endpoint, you must enable network policy support for the subnet that contains the private endpoint... Network policies aren't supported for private endpoints by default." This directly confirms that configuring the network policy is the required first action.
Source:
https://learn.microsoft.com/en-us/azure/private-link/disable-private-endpoint-network-policy

2. Microsoft Learn Azure Private Link Documentation. "What is a private endpoint? - Network security groups (NSG)." This section clarifies the relationship between NSGs and private endpoints: "To use an NSG with a private endpoint, you must enable network policy support for the subnet." This reinforces that enabling the policy is a prerequisite.
Source: https://learn.microsoft.com/en-us/azure/private-link/private-endpoint-overview#network-security-groups-nsg

Question: 17

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains a subnet named Subnet1. You plan to add a private endpoint to Subnet1. You need to ensure that you can route traffic between the private endpoint and the Azure Private Link service by using a user-defined route. What should you do first on Subnet1?

A. Enable network policy.

B. Enable delegation.

C. Create a service endpoint.

D. Provision a Standard Azure load balancer.

Answer:

A

Explanation:

By default, network policies such as User-Defined Routes (UDRs) and Network Security Groups (NSGs) are disabled for private endpoints on the subnet in which they are configured. This default behavior ensures that traffic to the Private Link service is not inadvertently blocked or misrouted. To allow a UDR to manage and route traffic originating from a private endpoint, you must first explicitly enable network policy enforcement on the subnet. This is achieved by setting the privateEndpointNetworkPolicies property for the subnet to Enabled. This is the prerequisite step on the subnet before a UDR can take effect on private endpoint traffic.

Why Incorrect Options are Wrong:

B. Enable delegation: Subnet delegation is used to dedicate a subnet to a specific Azure service (e.g., Azure NetApp Files), which is not required for private endpoint UDR functionality.
C. Create a service endpoint: Service endpoints are a different virtual network integration technology and do not influence how network policies are applied to private endpoints.
D. Provision a Standard Azure load balancer: A Standard Load Balancer is typically used on the service-provider side of a Private Link service, not on the client subnet to enable routing policies.

References:

1. Microsoft Learn Azure Private Link Documentation. "Manage network policies for private endpoints." This document states, "Network policies like Network security groups (NSG) and User Defined Routes (UDR) aren't supported for private endpoints. To use an NSG or UDR with a private endpoint, you must enable the network policy support for the private endpoint. This setting is configured on the subnet that contains the private endpoint."
Source:
https://learn.microsoft.com/en-us/azure/private-link/disable-private-endpoint-network-policy

(Section: "Enable network policy support for private endpoints")
2. Microsoft Learn Azure Private Link Documentation. "What is Azure Private Endpoint?" Under the "Network security of private endpoints" section, it clarifies, "By default, NSGs and UDRs don't apply to private endpoints. You can enable NSGs and UDRs for your private endpoints by enabling the privateEndpointNetworkPolicies property on the subnet that hosts the private endpoints."
Source: https://learn.microsoft.com/en-us/azure/private-link/private-endpoint-overview#network-security-of-private-endpoints

Question: 18

You have two Azure subscriptions named Sub1 and Sub2. Sub1 contains a virtual machine named VM1. You plan to make VM1 available to the resources in Sub2 by using Azure Private Link. You need to ensure that the private link service can be configured to provide access to VM1. What should you configure in Sub1 first?

A. a service endpoint

B. an Azure Private DNS zone

C. a private endpoint

D. an Azure load balancer

Answer:

D

Explanation:

To expose a service on a virtual machine (VM1) using Azure Private Link Service, the VM must be placed in the backend pool of a Standard SKU internal load balancer. The Private Link service is then created in the same virtual network (in Sub1) and points to the frontend IP configuration of this load balancer. This is a fundamental prerequisite for making the service available to consumers in other virtual networks or subscriptions. Therefore, configuring the Azure Load Balancer is the first required step in the provider subscription.

Why Incorrect Options are Wrong:

A. A service endpoint is used to secure Azure PaaS resources to a virtual network, not for exposing a VM-based service to another VNet.
B. An Azure Private DNS zone is used on the consumer side (Sub2) to resolve the private endpoint's name to its private IP address.
C. A private endpoint is the network interface created in the consumer's virtual network (in Sub2) to connect to the Private Link service.

References:

1. Microsoft Documentation Azure Private Link What is Azure Private Link service?
Reference: In the "Prerequisites" section, it explicitly states: "Your application runs behind a standard Azure Load Balancer in your virtual network." This confirms that the load balancer is a mandatory prerequisite.
Source: Microsoft Learn, "What is Azure Private Link service?", learn.microsoft.com/en-us/azure/private-link/private-link-service-overview#prerequisites.
2. Microsoft Documentation Azure Private Link Quickstart: Create a Private Link service by using the Azure portal

Reference: The "Prerequisites" section of this tutorial lists the creation of a virtual network, virtual machines, and a standard internal load balancer as the initial steps before creating the Private Link service itself.
Source: Microsoft Learn, "Quickstart: Create a Private Link service by using the Azure portal", learn.microsoft.com/en-us/azure/private-link/create-private-link-service-portal?tabs=internal#prerequisites.
3. Microsoft Documentation Azure Private Link Private endpoint overview
Reference: This document describes a private endpoint as "a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service that's powered by Azure Private Link." This clarifies that the private endpoint is the consumer-side component.
Source: Microsoft Learn, "What is a private endpoint?", learn.microsoft.com/en-us/azure/private-link/private-endpoint-overview.

Question: 19

You have two Azure virtual networks named VNet1 and VNet2 that are peered with each other. VNet1 hosts 10 virtual machines that contain web servers. VNet2 hosts five virtual machines that contain database servers. You need to configure a security solution that meets the following requirements: • Ensures that the database servers can accept connections only from the web servers • Ensures that the web servers can initiate connections only to the database servers • Ensures that all network security groups (NSGs) are associated only with subnets • Use application security groups to implement the solution What is the minimum number of application security groups required?

A. 1

B. 2

C. 4

D. 8

Answer:

B

Explanation:

The solution requires grouping virtual machines by their function to apply specific network security rules. One Application Security Group (ASG) is needed to group all the web servers in VNet1. A second ASG is needed to group all the database servers in VNet2.
With these two ASGs, you can create Network Security Group (NSG) rules that specifically allow traffic from the "web server" ASG to the "database server" ASG and deny other traffic. This is the minimum number of ASGs required to logically group the two distinct application tiers and enforce the required traffic flow.

Why Incorrect Options are Wrong:

A. 1: A single ASG cannot differentiate between web servers and database servers, making it impossible to create rules that restrict traffic between them.
C. 4: While you could create more ASGs, the question asks for the minimum number required. Two ASGs are sufficient to group the two distinct server roles.
D. 8: This is excessive. Two ASGs are sufficient to group the two distinct server roles and meet all the security requirements.

References:

1. Microsoft Learn, Azure Virtual Network Documentation, "Application security groups": "Application security groups enable you to configure network security as a natural extension of an application's structure, allowing you to group virtual machines and define network security policies based on those groups... In the security rule, you can specify an application security group as the source and destination." This directly supports the need for two groups: one for the source (web servers) and one for the destination (database servers).
2. Microsoft Learn, Azure Virtual Network Documentation, "Tutorial: Filter network traffic with a network security group using the Azure portal": In the section "Create application security groups," the tutorial demonstrates a similar scenario by creating two distinct ASGs: one for web servers (myAsgWebServers) and another for management servers (myAsgMgmtServers). This confirms the pattern of using one ASG per server role.
3. Microsoft Learn, Azure Virtual Network Documentation, "Network security groups overview": Under the "Application security groups" section, it states: "This capability allows you to use ASGs as the source and destination in a security rule." This reinforces that to define a rule between two distinct sets of servers, you need two corresponding ASGs.

Question: 20

HOTSPOT You have an Azure subscription. The subscription contains multiple Azure SQL Database resources and a virtual network named VNet1 that has five subnets. All the subnets are associated with a network security group (NSG) named NSG1. NSG1 blocks all outbound traffic, unless specifically allowed by a rule. Each subnet contains 50 virtual machines. Multiple virtual machines host instances of SQL Server on Virtual Machines and will be configured to replicate with the Azure SQL Database resources. You need to configure a new outbound rule in NSG1 to allow the SQL Server on Virtual Machines instances to connect to the Azure SQL Database resources. The solution must meet the following requirements: • Minimize modifications to NSG1 when additional instances of SQL Server on Virtual Machines are deployed. • Ensure that only SQL Server on Virtual Machines instances can connect to the Azure SQL Database resources. How should you configure each setting for the new outbound rule? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer:

Source: Application security group

Destination: Service Tag

Explanation:

To satisfy the requirements, the outbound Network Security Group (NSG) rule should be configured as follows:

• Source: Application security group (ASG) An ASG allows you to group virtual machines with similar functions, in this case, the SQL Server on Virtual Machines instances. By setting the ASG as the source, the rule applies only to these specific VMs. When new SQL Server VMs are deployed, you can simply add their network interfaces to this ASG. This approach fulfills both

requirements: it minimizes modifications to the NSG rule itself and ensures only the specified VMs can initiate outbound connections under this rule.

• Destination: Service Tag The destination is Azure SQL Database, which is an Azure PaaS service. A Service Tag represents a group of IP address prefixes for a specific Azure service. By using the Sql service tag as the destination, you allow traffic to all Azure SQL Database endpoints. Microsoft automatically manages and updates the IP addresses associated with the service tag, ensuring the rule remains current and effective without manual intervention. Using an IP address list would be brittle and violate the principle of minimizing modifications.

References:

Microsoft Azure Documentation, "Application security groups": "Application security groups enable you to configure network security as a natural extension of an application's structure... You can specify an application security group as the source and destination in a network security rule... This capability removes the need for manual maintenance of explicit IP addresses." This directly supports using an ASG to group the SQL Server VMs and minimize future changes.
Microsoft Azure Documentation, "Virtual network service tags": "A service tag represents a group of IP address prefixes from a given Azure service. Microsoft manages the address prefixes encompassed by the service tag and automatically updates the service tag as addresses change, minimizing the complexity of frequent updates to network security rules." This document confirms that using a service tag (e.g., Sql) is the appropriate method for targeting an Azure service like Azure SQL Database as a destination.

Question: 21

You have an Azure subscription that contains 100 network security groups (NSGs). You need to ensure that you log the application of specific NSG rules. Which type of log should you configure?

A. flow log

B. activity log

C. Azure resource log

D. audit log

Answer:

A

Explanation:

NSG flow logs are a feature of Azure Network Watcher designed specifically to log IP traffic flowing through a network security group (NSG). These logs capture data on a per-rule basis, recording details such as source/destination IP, port, protocol, and the traffic decision (Allowed or Denied). This directly addresses the need to log the application of specific NSG rules to actual network traffic, providing visibility into which rules are processing which flows.

Why Incorrect Options are Wrong:

B. activity log: This log records management plane operations, such as creating or modifying an NSG rule, not the data plane operation of a rule being applied to traffic.
C. Azure resource log: This is a general category of logs. While NSG flow logs are a type of resource log, "flow log" is the specific and correct feature for this task.
D. audit log: This is an older name for the Azure Activity Log and, like the activity log, it tracks management-level events, not the processing of network traffic.

References:

1. Microsoft Learn. (2023). Introduction to flow logging for network security groups. "A network security group flow log is a feature of Azure Network Watcher that allows you to log information about IP traffic flowing through a network security group... Flow logs are written in JSON format and show outbound and inbound flows on a per-rule basis..." This source confirms that flow logs operate on a per-rule basis.
2. Microsoft Learn. (2023). Tutorial: Log network traffic to and from a virtual machine using the Azure portal. In the "View flow log" section, the documentation shows an example of a flow log entry. The log format includes properties for the rule that was applied and the traffic decision ('A' for Allowed, 'D' for Denied), demonstrating how it logs the application of rules.
3. Microsoft Learn. (2023). Azure platform logs overview. This document distinguishes between different log types. It clarifies that the Activity log contains entries for "control-plane events," while

Resource logs (which include NSG flow logs) contain information about "operations that were performed by an Azure resource" (the data plane), confirming that the Activity/Audit log is incorrect for this scenario.

Question: 22

You are planning an Azure deployment that will contain three virtual networks in the East US Azure region as shown in the following table.

Name	Description
Vnet1	Hub virtual network for shared services
Vnet2	Virtual machines for the IT department
Vnet3	Virtual machines for the research department

A Site-to-Site VPN will connect Vnet1 to your company's on-premises network. You need to recommend a solution that ensures that the virtual machines on all the virtual networks can communicate with the on-premises network- The solution must minimize costs. What should you recommend for Vnet2 and Vnet3?

A. service endpoints

B. route tables

C. VNet-to-VNet VPN connections

D. peering

Answer:

D

Explanation:

Virtual network (VNet) peering is the most appropriate and cost-effective solution. By peering VNet2 and VNet3 with VNet1, you can configure gateway transit. This feature allows the peered VNets (VNet2 and VNet3) to use the VPN gateway in VNet1 to access the on-premises network. This creates a hub-and-spoke network topology where VNet1 acts as the hub. This approach centralizes the connection point and avoids the significant cost and management overhead of deploying separate VPN gateways in VNet2 and VNet3. Since all VNets are in the same region, there are no data transfer charges for the peering connection itself.

Why Incorrect Options are Wrong:

A. service endpoints: These are used to secure Azure service resources to a virtual network, not for connecting VNets to each other or to on-premises networks.
B. route tables: While route tables are used to direct traffic to the gateway once peering is established, they do not create the actual connectivity path between the VNets.
C. VNet-to-VNet VPN connections: This would require deploying a VPN gateway in each VNet, incurring substantial additional hourly costs and management complexity, violating the cost minimization requirement.

References:

1. Microsoft Learn Azure Virtual Network peering: Under the "Gateways and on-premises connectivity" section, the documentation states, "With virtual network peering, you can configure gateway transit from a virtual network that contains a VPN gateway to provide on-premises connectivity to the peered virtual network." This directly supports using peering with gateway transit for the described scenario.
Source: https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-peering-overview#gateways-and-on-premises-connectivity
2. Microsoft Learn Configure VPN gateway transit for virtual network peering: This tutorial explicitly details the hub-and-spoke model. The "About gateway transit" section explains, "Gateway transit allows peered virtual networks to share the gateway... This saves cost and reduces management overhead because you only need to deploy and manage one VPN gateway for all the peered virtual networks."
Source: https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-peering-gateway-transit#about-gateway-transit
3. Azure Pricing Virtual Network: The official pricing page confirms the cost-effectiveness. Under the "VNet Peering" section, it shows that for "Data transfer between virtual networks in the same region," both Ingress and Egress are free. This contrasts with the costs associated with deploying multiple VPN gateways.
Source: https://azure.microsoft.com/en-us/pricing/details/virtual-network/

Question: 23

HOTSPOT You have an Azure subscription that contains a virtual machine named VM1 and a virtual network named Vnet1. Vnet1 contains three subnets named Subnet1, Subnet2 and GatewaySubnet. VM1 is connected to Subnet 1. You plan to deploy a new virtual machine named VM2 that will perform network traffic routing and inspection. You need to ensure that all the traffic from VM1 to the internet will be routed through VM2. What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer:

Deploy VM2 to: Subnet2

Create a custom route table and associate the table with: Subnet1

Explanation:

To route all internet-bound traffic from VM1 through a new virtual machine (VM2), which acts as a Network Virtual Appliance (NVA), you must configure a User-Defined Route (UDR).

• Deploy VM2 to Subnet2: It is a best practice to deploy the NVA (VM2) in its own dedicated subnet. This isolates the appliance for easier management and security rule application. The GatewaySubnet is reserved exclusively for Azure's VPN and ExpressRoute gateways and cannot be used for deploying virtual machines. Placing VM2 in Subnet1 alongside VM1 is possible but not recommended for proper network segmentation.

• Associate Route Table with Subnet1: The custom route table must be associated with the subnet containing the source of the traffic you wish to redirect. Since VM1 is in Subnet1, the route table must be applied to Subnet1. This route table will contain a rule that directs all traffic destined for the internet (0.0.0.0/0) to the private IP address of VM2 as the next hop, effectively forcing traffic through the NVA for inspection.

References:

Microsoft Azure Documentation, Route network traffic with a route table using the Azure portal.
This tutorial demonstrates the exact scenario. It states, "You create a route table and associate the route table to a subnet. The route directs traffic destined for the public subnet through the NVA virtual machine." The architecture shown places the NVA in a separate subnet and associates the route table with the source subnet.
Microsoft Azure Documentation, Virtual network traffic routing.
Under the "User-defined" routing section, it explains: "You can create custom, or user-defined(static), routes in Azure to override Azure's default system routes... You associate a route table to zero or more virtual network subnets." This confirms that to control traffic from a subnet, you must associate the route table with that specific subnet.
Microsoft Azure Documentation, What is Azure Virtual Network Gateway?.
In the overview, it specifies the requirement for a special subnet: "A virtual network gateway requires a specific subnet called the gateway subnet. The gateway subnet is part of the virtual network IP address range that you specify when you configure your virtual network. It contains the IP addresses that the virtual network gateway resources and services use." This confirms the dedicated purpose of GatewaySubnet, precluding VM deployment.

Question: 24

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
storage1	Storage account	*None*
storage2	Storage account	*None*
DB1	Azure SQL Database	*None*
VNet1	Virtual network	Peered with VNet2 Contains two subnets that each contains 10 virtual machines
VNet2	Virtual network	Peered with VNet1 Contains two subnets that each contains 10 virtual machines

You need to ensure that the virtual machines can access storage1, storage2, and DB1 by using service endpoints. What is the minimum number of service endpoints you should create?

A. 2

B. 3

C. 4

D. 12

Answer:

B

Explanation:

A service endpoint is configured on a virtual network subnet for a specific Azure service type. The minimum number of endpoints is determined by the number of subnets and the types of services they need to access, considering the scope of each endpoint.

1. VNet1 (East US): To allow VM1 to access the required resources, its subnet needs two endpoints:

One for Microsoft.Storage to access both storage1 and storage2. Storage service endpoints are global.

One for Microsoft.Sql to access DB1. This works because both VNet1 and DB1 are in the same region (East US).

2. VNet2 (West US): To allow VM2 to access the required resources, its subnet needs one endpoint:

One for Microsoft.Storage to access both storage1 and storage2.

A Microsoft.Sql service endpoint cannot be created on VNet2 to access DB1 because SQL

service endpoints are regional and only work for SQL servers within the same region as the VNet. Therefore, only 3 endpoints can be created to fulfill the possible requirements.

Why Incorrect Options are Wrong:

A. 2: This is incorrect. It only accounts for the two endpoints required for VNet1 and neglects the storage access needed from VNet2.
C. 4: This would be correct if SQL service endpoints were global like storage endpoints. However, they are regional, making the VNet2-to-DB1 connection via service endpoint impossible.
D. 12: This is incorrect. The number of endpoints is not a product of the number of VMs and target resources; it's based on service types per source subnet.

References:

1. Microsoft Documentation Virtual Network service endpoints: Under the "Limitations" section, it states, "For Azure SQL Database, a service endpoint applies only to Azure SQL traffic in a virtual network's region. For Azure Storage, endpoints extend to all regions." This directly supports the calculation of 3 endpoints.
Source: Microsoft Learn, "Virtual Network service endpoints", Section: "Limitations".
2. Microsoft Documentation Configure virtual network service endpoints: This document explains that service endpoints are enabled on a subnet for a specific service. "You can enable a service endpoint for supported Azure services within a virtual network. The endpoint is enabled on a subnet that you configure in a virtual network." This confirms that endpoints are configured per-subnet, per-service.
Source: Microsoft Learn, "Tutorial: Restrict network access to PaaS resources with virtual network service endpoints", Section: "Enable a service endpoint".

Question: 25

DRAG DROP

You have two Azure subscriptions named Sub1 and Sub2 that contain the resources shown in the following table.

[image could not be rendered]

VNet1 and VNet2 are NOT connected.

You plan to create an Azure Private Link service named Link1 that will be used to connect VNet1 and VNet2.

You need to ensure that Link1 meets the following requirements:

• Ensures that VM1 can connect only to a web app hosted on VM2

• Prevents VM1 from connecting to the other resources that are connected to VNet2

Which additional resources should you create for each virtual network? To answer, drag the appropriate resources to the correct virtual networks. Each resource may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

[image could not be rendered]

Answer:

VNet1: A private endpoint

VNet2: A load balancer

Explanation:

To establish the required secure and isolated connection using Azure Private Link, you need two key components configured in the respective virtual networks.

• VNet2 (Service Provider side): To expose the web app on VM2 as a Private Link service, the virtual machine must be placed in the backend pool of a Standard SKU load balancer. The Private Link service is then associated with the frontend IP configuration of this load balancer. This setup makes the specific service on VM2 available for private connection without exposing any other

resources (like VM3 or VM4) or the VNet itself.

• VNet1 (Consumer side): To consume the service exposed from VNet2, you must create a private endpoint within VNet1. This private endpoint creates a network interface in VNet1 with a private IP address from VNet1's address space. When VM1 sends traffic to this private IP, it's securely and privately routed over the Microsoft backbone network directly to the service on VM2, completely bypassing the public internet and preventing access to other resources in VNet2.

References:

Microsoft Learn, "Azure Private Link service concepts": This document explicitly states the prerequisites for a Private Link service. In the "Private Link service" section, it says, "Your service running behind a standard load balancer can be enabled for Private Link access." This confirms the requirement for a load balancer in VNet2.
Microsoft Learn, "What is a private endpoint?": This resource defines the role of a private endpoint. It states, "A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service that's powered by Azure Private Link." This validates the need for a private endpoint in VNet1 to consume the service.
Microsoft Learn, "Quickstart: Create a Private Link service by using an ARM template": This tutorial provides a step-by-step guide for deployment. The architecture diagram and template details show a Standard Load Balancer deployed in the service provider's virtual network and a Private Endpoint in the consumer's virtual network to establish the connection.

Question: 26

HOTSPOT You have an Azure subscription that contains a virtual network named VNet1. VNet1 uses an IP address space of 192.168.0.0/24. You plan to deploy Azure virtual machines and Azure Bastion to VNet1. You need to recommend an IP subnetting configuration for VNet1. The solution must maximize the number of IP addresses that can be assigned to the virtual machines What should you recommend? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

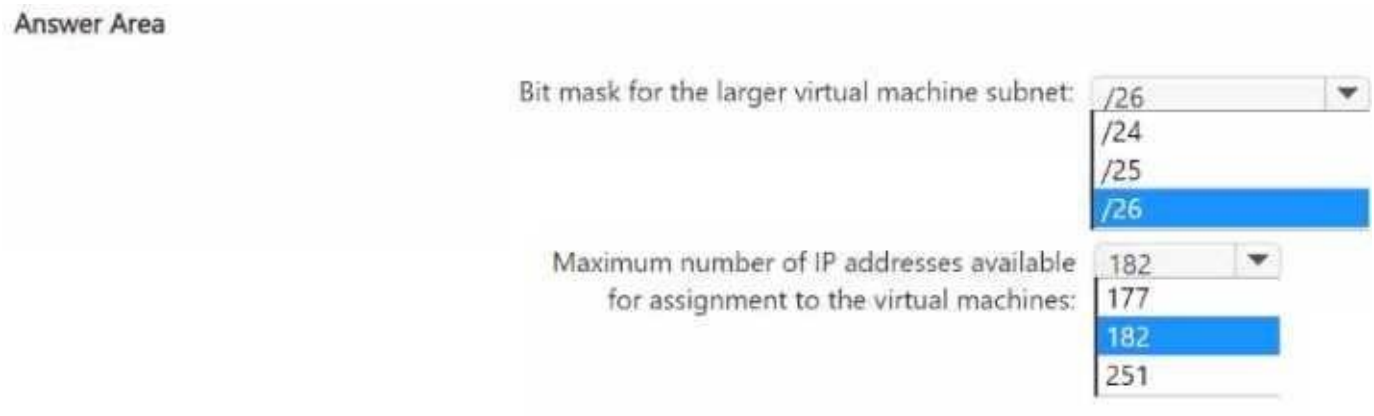

Answer:

Bit mask for the larger virtual machine subnet: /25

Maximum number of IP addresses available for assignment to the virtual machines: 182

Explanation:

The problem requires subnetting a 192.168.0.0/24 virtual network to accommodate Azure Bastion and virtual machines, while maximizing the number of available IPs for the VMs.

• Azure Bastion Subnet Requirement: Azure Bastion requires a dedicated subnet named AzureBastionSubnet with a minimum size of /26. To maximize the remaining space for VMs, we should allocate this minimum size to the Bastion subnet. A /26 subnet contains 2(3226)=26=64 total IP addresses.

• VM Subnet Allocation: The total VNet space is /24, which has 28=256 addresses. After allocating a /26 subnet (64 addresses) for Bastion, 256 - 64 = 192 addresses remain for the VMs. This remaining space cannot be represented by a single CIDR block. To utilize all 192 addresses, it must be divided into a /25 subnet (128 addresses) and a /26 subnet (64 addresses).

• Answering the Questions:

• Bit Mask: The two subnets created for the VMs are /25 and /26. The question asks for the bit mask of the larger subnet, which is /25.

• Available IPs: Azure reserves 5 IP addresses in every subnet. The total number of available IPs for VMs is the sum of available IPs from both VM subnets:

• For the /25 subnet: 128 total IPs - 5 reserved IPs = 123 available IPs.

• For the /26 subnet: 64 total IPs - 5 reserved IPs = 59 available IPs.

• Total available for VMs = 123 + 59 = 182.

References:

Azure Bastion Subnet: According to the official Microsoft Azure documentation, "When you deploy Azure Bastion using any SKU except the Developer SKU, Bastion requires a dedicated subnet named AzureBastionSubnet. You must create this subnet in the same virtual network that you want to deploy Bastion to. The subnet must have a minimum /26 prefix or larger."
Source: Microsoft Docs, Create a bastion host. (This information is found in the "Prerequisites" and "AzureBastionSubnet" sections of the documentation for creating a Bastion host).
Azure VNet IP Address Reservation: In each Azure subnet, the first four IP addresses and the last IP address are reserved for protocol conformance, totaling 5 reserved addresses per subnet.
Source: Microsoft Docs, Azure Virtual Network frequently asked questions (FAQ), "Are there any restrictions on using IP addresses within these subnets?".

Question: 27

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains an Azure Virtual Desktop host pool named Pool1. You need to implement Azure Firewall and TLS inspection for all the outbound traffic from Pool1. Which two resources should you configure? Each correct answer present part of the solution. NOTE: Each correct answer is worth one point

A. an Azure Private DNS zone

B. a private endpoint

C. an Azure key vault

D. an Azure NAT gateway

E. a Microsoft Entra enterprise app

F. a managed identity

Answer:

C, F

Explanation:

To implement TLS inspection with Azure Firewall Premium, two primary resources are essential for certificate management and authentication. First, an Azure Key Vault is required to securely store the intermediate Certificate Authority (CA) certificate and its corresponding private key. The firewall uses this certificate to generate certificates on-the-fly for inspected traffic. Second, a Managed Identity (specifically, a user-assigned one) must be associated with the firewall. This identity is granted permissions to the Azure Key Vault, allowing the firewall to securely access the certificate without exposing credentials. These two components are fundamental prerequisites for enabling the TLS inspection feature.

Why Incorrect Options are Wrong:

A. an Azure Private DNS zone: This is used for custom DNS resolution within a virtual network and is not a direct requirement for enabling TLS inspection.
B. a private endpoint: This provides private, inbound connectivity to Azure services and is not used for inspecting outbound traffic from a host pool.
D. an Azure NAT gateway: This provides Source Network Address Translation (SNAT) for outbound connectivity, a function that Azure Firewall can perform natively. It is not a prerequisite for TLS inspection.
E. a Microsoft Entra enterprise app: This is an identity construct for applications (for SSO, etc.) and is not used to grant a network appliance like Azure Firewall access to certificates.

References:

1. Microsoft Documentation - Azure Firewall Premium features: Under the "TLS inspection" section, it states: "To properly configure TLS Inspection, you must create a Firewall Policy, enable TLS inspection, and provide a Key Vault certificate and a Managed Identity."
Source: Microsoft Learn, Azure Firewall Premium features, Section: "TLS inspection".
2. Microsoft Tutorial - Deploy and configure Azure Firewall Premium: The tutorial prerequisites and configuration steps explicitly detail the creation and use of both a Key Vault to hold the certificate and a user-assigned managed identity to provide the firewall with access to that Key Vault.
Source: Microsoft Learn, Tutorial: Deploy and configure Azure Firewall Premium, Sections: "Prerequisites" and "Configure the firewall policy".
3. Microsoft Documentation - Azure Firewall certificates: This document clarifies the certificate requirements for TLS inspection, stating, "The intermediate CA certificate must be stored in an Azure Key Vault... The firewall uses a Managed Identity associated with it to get the certificate from Key Vault."
Source: Microsoft Learn, Azure Firewall Premium certificates, Section: "Certificate requirements".

Question: 28

DRAG DROP

You have two on-premises datacenters.

You have an Azure subscription that contains four virtual networks named VNet1, VNet2, VNet3, and VNet4.

You create an Azure virtual WAN named VWAN1. VWAN1 contains a single virtual hub that is connected to both on-premises datacenters and all the virtual networks in a full mesh topology.

You create a route table named RT1.

You need to configure VWAN1 to meet the following requirements:

• Connectivity between VNet1 and VNet2 and both on-premises datacenters must be allowed.

• Connectivity between VNet3 and VNet4 and both on-premises datacenters must be allowed.

• VNet1 and VNet2 must be isolated from VNet3 and VNet4.

How should you configure routing for VNet1 and VNet2 and for both on-premises datacenters? To answer, drag the appropriate route tables and route table propagation to the correct requirements. Each route table and route table propagation may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

[image could not be rendered]

Answer:

VNet1 and VNet2:

• Associated route table: RT1;

• Propagating to route tables: RT1 and Default

On-premises datacenters:

• Associated route table: Default

• Propagating to route tables: RT1 and Default

Explanation:

This configuration establishes two distinct routing domains to satisfy the isolation requirement, using the Default and RT1 route tables.

• On-premises Datacenters:

• Propagating to RT1 and Default: The on-premises connection must advertise its routes to both routing domains so that all virtual networks (VNet1/VNet2 and VNet3/VNet4) can reach it.

• Associated route table Default: This configuration effectively turns the Default table into a "superset" routing table. It will learn routes from all connections that propagate to it (VNet1/VNet2, VNet3/VNet4, and on-premises). By associating with this table, the on-premises datacenter can learn the routes to all four VNets, enabling full connectivity.

• VNet1 and VNet2:

• Associated route table RT1: By associating VNet1 and VNet2 with the custom RT1 table, their outbound traffic is governed by this table, placing them in an isolated routing domain.

• Propagating to RT1 and Default:

• They propagate to RT1 so they can learn each other's routes, enabling connectivity within their group. The RT1 table will only contain routes for VNet1, VNet2, and the on-premises network (which also propagates to RT1), thus isolating them from VNet3 and VNet4.

• They propagate to Default so that the on-premises datacenter (which is associated with Default) can learn their routes and establish connectivity to them.

This setup correctly allows VNet1 and VNet2 to communicate with each other and with the on-premises datacenters while preventing them from initiating traffic to VNet3 and VNet4.

References:

Microsoft Azure Documentation: Virtual WAN routing scenarios: This document details various routing configurations. The "Isolating VNets" and "Shared services VNet" scenarios are particularly relevant. The principle is to use separate route tables for different groups of VNets. A shared resource (like the on-premises connection here) must propagate its routes to all route tables to be reachable, and it must associate with a table that learns routes from all VNet groups

to achieve full outbound connectivity.
Reference: Microsoft Docs, "About virtual hub routing," section on "Association" and "Propagation."
Microsoft Azure Documentation: How to configure virtual hub routing: This guide provides step-by-step instructions that align with the logic used in the solution. It explains that association dictates the routes a connection learns (for egress traffic), while propagation advertises a connection's prefixes to be learned by others (for ingress traffic).
Reference: Microsoft Docs, "How to configure virtual hub routing," sections on creating route tables, association, and propagation.

Question: 29

DRAG DROP

Your on-premises network uses an IP address space of 10.0.0.0/20.

You have an Azure subscription that contains the resources shown in the following table.

[image could not be rendered]

The on-premises network is connected to HubVnet by using a Site-to-Site (S2S) VPN.

You deploy an Azure firewall named AZFW1 to HubVNet.

You need to ensure that AZFW/1 can inspect all the traffic between the on-premises network and SpokeVNet.

What should you do in RT1? To answer, drag the appropriate destination to the correct route. Each resource may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

[image could not be rendered]

Answer:

Add a route for 10.0.0.0/20 and specify AZFW1 as the next hop for: All the subnets on SpokeVNet

Add a route for 192.168.0.0/20 and specify AZFW1 as the next hop for: GatewaySubnet on HubVNet

Explanation:

This configuration establishes a secure hub-and-spoke topology where Azure Firewall (AZFW1) inspects traffic between different networks.

• Spoke to On-Premises Traffic: To inspect traffic originating from SpokeVNet (192.168.0.0/20) and destined for the on-premises network (10.0.0.0/20), a User-Defined Route (UDR) is required. This route must be associated with the subnets within SpokeVNet. It directs any traffic intended for the on-premises address space to the Azure Firewall for inspection before it's sent to the hub's VPN gateway.

• On-Premises to Spoke Traffic: To inspect traffic coming from the on-premises network and

destined for SpokeVNet, a UDR must be applied to the GatewaySubnet within HubVNet. This route intercepts the inbound traffic as it arrives at the VPN gateway and forwards it to the Azure Firewall before it can proceed to the peered SpokeVNet.

References:

Microsoft Learn, Tutorial: Deploy and configure Azure Firewall in a hybrid network using the Azure portal. This official tutorial details the exact steps for this scenario.
Under the section "Create routes," it specifies creating a route for the spoke network on the GatewaySubnet's route table. It states, "This route sends any traffic from the gateway subnet to the Spoke-VNet through the firewall."
It also describes creating a route for the on-premises network and associating it with the spoke subnet's route table to ensure return traffic is also inspected.
Microsoft Learn, Hub-spoke network topology in Azure. This document outlines the architecture pattern.
Under the section "Spoke connectivity," it explains that User-Defined Routes (UDRs) are necessary to "make sure traffic from spoke virtual networks transits through the central hub" and can be directed to virtual appliances like Azure Firewall.
Microsoft Learn, Azure virtual network traffic routing. This document explains how Azure routes traffic.
Under the section "User-defined," it clarifies that you can create custom route tables with specific routes (UDRs) that override Azure's default routing. This is the mechanism used to force traffic through the firewall appliance. The documentation notes that to route traffic to a virtual appliance, you select "Virtual appliance" as the next hop type and specify the appliance's IP address.

Question: 30

DRAG DROP

You have an Azure Web Application Firewall (WAF) v2 tier named AG1 on an Azure application gateway. AG1 has a policy named Policy1.

You need to add a custom rule to Policy1. The rule must block all requests from IP addresses in a specific IP address range.

Which four PowerShell cmdlets should you run in sequence? To answer, move the appropriate cmdlets from the list of cmdlets to the answer area and arrange them in the correct order.

[image could not be rendered]

Answer:

1: New-AzApplicationGatewayFirewallMatchVariable

2: New-AzApplicationGatewayFirewallCondition

3: New-AzApplicationGatewayFirewallCustomRule

4: Set-AzApplicationGatewayFirewallPolicy

Explanation:

To create a custom rule in Azure Web Application Firewall (WAF), you must follow a specific sequence. First, you define the part of the request to inspect, which in this case is the source IP address (RemoteAddr), using New-AzApplicationGatewayFirewallMatchVariable. Next, you create a condition with New-AzApplicationGatewayFirewallCondition to specify the operator (e.g., IPMatch) and the IP range to match against the variable. Then, these components are bundled into a rule that defines an action (Block) and priority using New-AzApplicationGatewayFirewallCustomRule. Finally, the newly created rule is applied to the existing WAF policy using Set-AzApplicationGatewayFirewallPolicy to enforce it.

References:

Source: Microsoft Learn, Official Documentation
Title: Create and use Web Application Firewall v2 custom rules on Application Gateway
Section: In the "IP address restriction example" section, the document outlines the precise order of PowerShell cmdlets required. It demonstrates creating a match variable for RemoteAddr, followed by a condition using IPMatch, then creating a custom rule with a "Block" action, and finally updating the policy. This directly corresponds to the correct answer sequence.

Step 1: $matchVariable = New-AzApplicationGatewayFirewallMatchVariable ...

Step 2: $condition = New-AzApplicationGatewayFirewallCondition ...

Step 3: $rule = New-AzApplicationGatewayFirewallCustomRule ...

Step 4: Set-AzApplicationGatewayFirewallPolicy -InputObject $policy ...

Question: 31

You have an Azure subscription that contains an instance of Azure Firewall Standard named AzFW1. You plan to enable the following: • TLS inspection • Threat intelligence • A network intrusion detection and prevention system (IDPS) What can you enable by using AzFW1?

A. TLS inspection only

B. threat intelligence only

C. TLS inspection and the IDPS only

D. threat intelligence and the IDPS only

E. TLS inspection, threat intelligence, and the IDPS

Answer:

B

Explanation:

Azure Firewall is available in multiple SKUs (Standard, Premium, and Basic), each offering a different set of capabilities. The instance in the scenario, AzFW1, is Azure Firewall Standard. Azure Firewall Standard includes threat intelligence-based filtering, which can be configured to alert on or deny traffic from known malicious IP addresses, FQDNs, and URLs. However, TLS inspection and the network intrusion detection and prevention system (IDPS) are advanced features exclusive to the Azure Firewall Premium SKU. Therefore, of the three features listed, only threat intelligence can be enabled on the existing AzFW1 Standard instance.

Why Incorrect Options are Wrong:

A. TLS inspection is an exclusive feature of Azure Firewall Premium and is not available in the Standard SKU.
C. Both TLS inspection and IDPS are exclusive features of Azure Firewall Premium, not the Standard SKU.
D. The IDPS is an exclusive feature of Azure Firewall Premium and cannot be enabled on the Standard SKU.
E. TLS inspection and IDPS are exclusive features of Azure Firewall Premium. Only threat intelligence is available on the Standard SKU.

References:

1. Microsoft Learn, Azure Firewall features. The "Feature comparison" table explicitly shows that "Threat intelligence" is supported by Azure Firewall Standard, while "TLS Inspection" and "IDPS" are only supported by Azure Firewall Premium.
Reference: Microsoft Learn. (2023). Azure Firewall features. Retrieved from

https://learn.microsoft.com/en-us/azure/firewall/features, under the "Feature comparison" table.
2. Microsoft Learn, Azure Firewall Premium features. This document details the capabilities exclusive to the Premium SKU, confirming that IDPS and TLS inspection are part of this tier.
Reference: Microsoft Learn. (2023). Azure Firewall Premium features. Retrieved from https://learn.microsoft.com/en-us/azure/firewall/premium-features, under the "IDPS" and "TLS inspection" sections.
3. Microsoft Learn, Azure Firewall threat intelligence-based filtering. This document confirms that threat intelligence is a core feature available in Azure Firewall Standard.
Reference: Microsoft Learn. (2023). Azure Firewall threat intelligence-based filtering. Retrieved from https://learn.microsoft.com/en-us/azure/firewall/threat-intel, under the "How it works" section.

Question: 32

You have an on-premises DNS server named Server1 that hosts a primary DNS zone named fabrikam.com. You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
VNet1	Virtual network	In the US West Azure region
VNet2	Virtual network	In the US West Azure region
VNet3	Virtual network	In the West Europe Azure region
VNet4	Virtual network	In the West Europe Azure region
contoso.com	Azure Private DNS zone	In the West Europe Azure region and linked to VNet1, VNet2, VNet3, and VNet4

Users on the on-premises network access resources on all the virtual networks by using a Site-to-Site (S2S) VPN. You need to deploy an Azure DNS Private Resolver solution that meets the following requirements: • Resources connected to the virtual networks must be able to resolve DNS names for fabrikam.com. • Server1 must be able to resolve the DNS names of the resources in contoso.com. • The solution must minimize costs and administrative effort. What is the minimum number of resolvers you should deploy?

A. 1

B. 2

C. 3

D. 4

Answer:

A

Explanation:

The scenario requires bidirectional DNS name resolution between an on-premises network and Azure Virtual Networks. An Azure DNS Private Resolver is the appropriate service for this. A single Azure DNS Private Resolver resource is designed to handle both inbound and outbound DNS queries.

To meet the requirements, you configure this single resolver with:

1. An inbound endpoint to receive DNS queries from the on-premises server (Server1) for contoso.com.
2. An outbound endpoint and a DNS forwarding ruleset to forward queries from Azure resources for fabrikam.com to the on-premises server.

Since both the required inbound and outbound endpoints can be created within a single Azure

DNS Private Resolver resource, only one resolver needs to be deployed, minimizing both cost and administrative effort.

Why Incorrect Options are Wrong:

B. 2: This is incorrect. While two endpoints (one inbound, one outbound) are needed, they are components of a single DNS Private Resolver resource. Deploying two separate resolver resources is unnecessary and increases costs.
C. 3: This is incorrect. There is no scenario described that would require three separate resolver resources. This would unnecessarily increase complexity and cost.
D. 4: This is incorrect. Deploying four resolvers is excessive and would violate the requirement to minimize costs and administrative effort.

References:

1. Microsoft Documentation - What is Azure DNS Private Resolver?: "A DNS private resolver is a service that enables you to query Azure DNS private zones from an on-premises environment and vice versa... A single private resolver can satisfy the DNS query requirements for a complex network scenario." This confirms that one resolver resource can handle the entire hybrid scenario.
2. Microsoft Documentation - Azure DNS Private Resolver endpoints and rulesets: "Inbound and outbound endpoints can be provisioned in the same private resolver... A private resolver can have many inbound and outbound endpoints, but only one of each is required for a hybrid DNS scenario." This explicitly states that one resolver resource can contain the necessary inbound and outbound components.
3. Microsoft Documentation - Azure DNS Private Resolver architecture: The architecture diagrams clearly show a single "DNS Private Resolver" resource containing both an "Inbound Endpoint" and an "Outbound Endpoint" to facilitate hybrid name resolution. This visual representation supports the deployment of a single resource.

Question: 33

HOTSPOT You have two Azure subscriptions. You need to perform the following actions in the East US Azure region of each subscription: • Deploy 50 virtual machines to availability zone 1. • Deploy 50 virtual machines to availability zone 2. • Deploy 50 virtual machines to availability zone 3. What is the minimum number of virtual networks and /25 subnets you should create? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

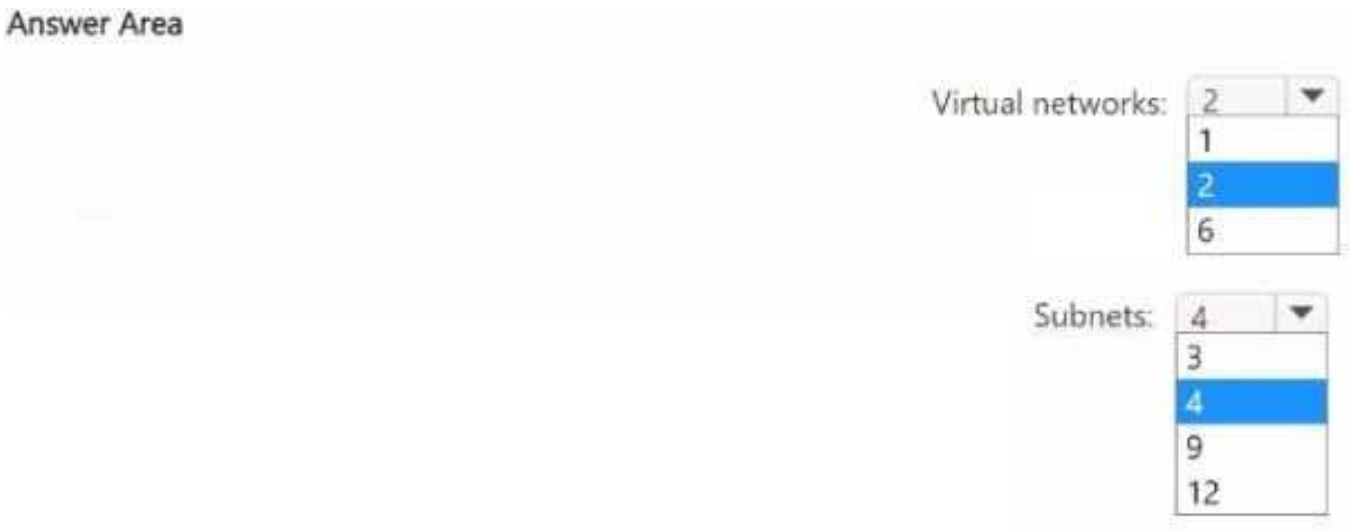

Answer:

Minimum number of virtual networks: 2
Minimum number of /25 subnets: 4

Explanation:

A virtual network is scoped to a single subscription and a single region. Because the deployment spans two subscriptions, a minimum of two virtual networks is required-one in each subscription. Each subscription must host a total of 150 virtual machines (50 VMs 3 availability zones). A /25 subnet provides 128 total IP addresses (2(32-25)). Azure reserves five IP addresses in every subnet, leaving 123 usable addresses. One /25 subnet is insufficient for the 150 VMs required in each subscription. Therefore, a minimum of two /25 subnets are needed per subscription to provide enough IP addresses (2 123 = 246 available IPs).
Across two subscriptions, the total minimum is four /25 subnets (2 subnets per subscription 2 subscriptions).

References:

1. Microsoft Learn. (2024). Azure Virtual Network concepts and best practices. Section: Scope. "A VNet is scoped to a subscription." and "An Azure VNet is scoped to a single region/location".
2. Microsoft Learn. (2024). Virtual Network FAQ. Section: How many IP addresses does Azure

reserve within each subnet?. "Azure reserves 5 IP addresses within each subnet. These are x.x.x.0-x.x.x.3 and the last address of the subnet."
3. Microsoft Learn. (2024). What is Azure Virtual Network?. Section: Virtual networks and subnets. "A subnet is a range of IP addresses in the VNet... A subnet spans all Availability Zones in a region." This confirms a single subnet can serve VMs across multiple zones, making the IP count the primary constraint.

Question: 34

HOTSPOT You have an on-premises network that includes the sites shown in the following table.

Site	Site Address space	Firewall private IP	Firewall public IP address
Paris	172.16.0.0/24	172.16.0.1	131.107.50.60
Amsterdam	172.16.1.0/24	172.16.1.1	131.107.70.80
Berlin	172.16.2.0/24	172.16.2.1	131.107.90.100

Each site is connected to the Internet by a firewall. All sites are connected to an SD-WAN. Each site is configured to propagate routes by using BGP. You have an Azure subscription that includes a virtual network named Vnet1 that contains a Virtual Network Gateway named Gateway 1. You create a local network gateway with the configuration shown in the gateway exhibit (Click the

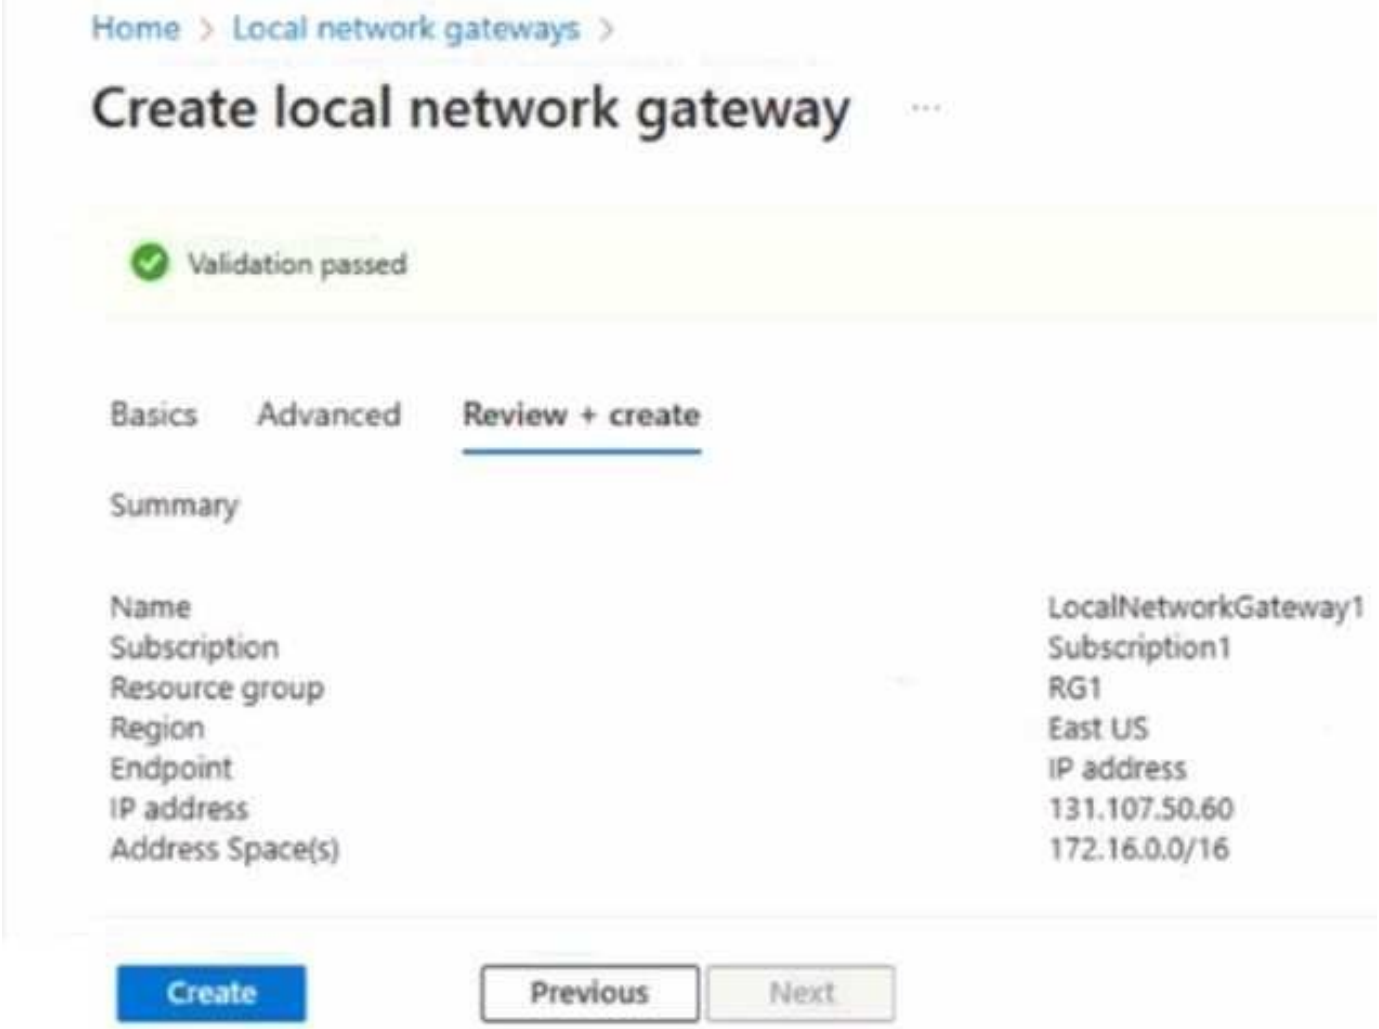

You create a Site-to-Site (S2S) connection with the configuration shown in connection exhibit. (Click the Connection tab)

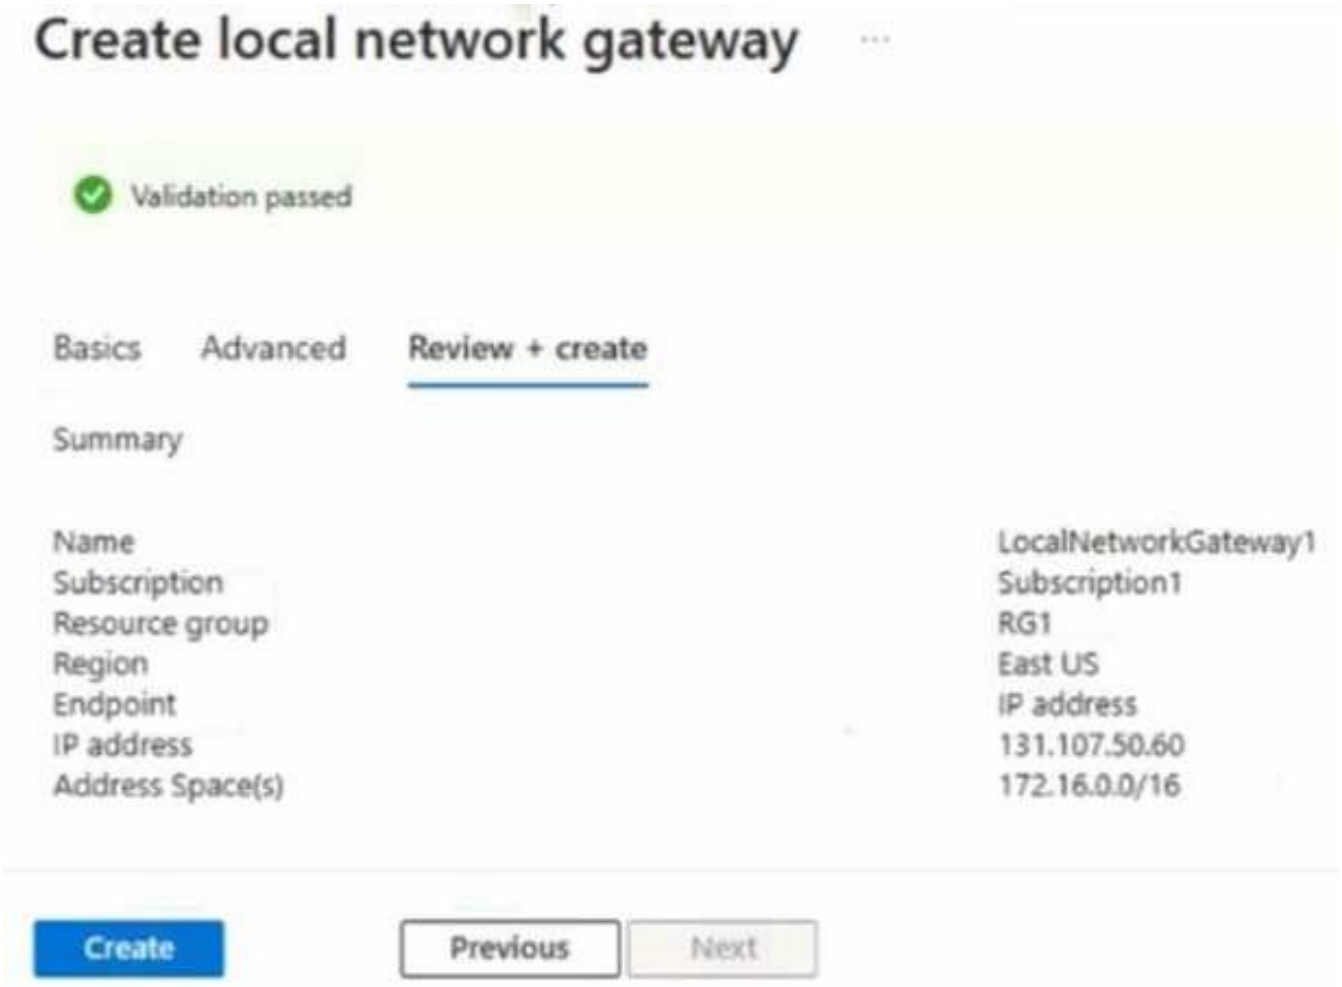

For each of the following statements, select Yes if the statement is true Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Users in the Berlin site can connect to resources in Vnet1 via VPN1.	○	○
To create a direct Site-to-Site connection to the Berlin site an additional Local Network Gateway is required.	○	○
To enable users in the Paris site to connect to Vnet1, the IP address of LocalNetworkGateway1 must be changed to 172.16.0.1.	○	○

Answer:

Yes

Yes

No

Explanation:

Users in the Berlin site can connect to resources in Vnet1 via VPN1: Yes The LocalNetworkGateway1 is configured with the address space 172.16.0.0/16, which is a summary route that includes the address spaces for Paris (172.16.0.0/24), Amsterdam (172.16.1.0/24), and Berlin (172.16.2.0/24). The VPN tunnel terminates at the Paris firewall (131.107.50.60). Since all on-premises sites are connected via an SD-WAN with BGP routing, traffic from VNet1 destined for Berlin will be sent over the VPN to the Paris firewall, which will then route it to the Berlin site. This is known as transitive routing.

To create a direct Site-to-Site connection to the Berlin site an additional Local Network Gateway is required: Yes In Azure, a Local Network Gateway (LNG) represents a specific on-premises site. It defines the public IP address of the on-premises VPN device and the address prefixes for that site. To establish a direct connection to the Berlin site, a new LNG must be created that specifies the Berlin firewall's public IP address (131.107.90.100) and its local address space (172.16.2.0/24). The existing LNG points to the Paris site's firewall.

To enable users in the Paris site to connect to Vnet1, the IP address of LocalNetworkGateway1 must be changed to 172.16.0.1: No The IP address field in a Local Network Gateway configuration requires the public IP address of the on-premises VPN device, which is used to establish the IKE/IPsec tunnel over the internet. The address 131.107.50.60 is the correct public IP for the Paris firewall. The address 172.16.0.1 is the firewall's private IP address, which is not reachable from the public internet and cannot be used to establish the VPN connection from Azure. Changing to the private IP would break the existing connection.

References:

Microsoft Documentation Azure VPN Gateway: "What is a local network gateway in Azure?" This document specifies that the local network gateway object contains settings such as the public IP address of the on-premises VPN device and the address prefixes that are on the on-premises network.
Microsoft Documentation Azure VPN Gateway: "Create a Site-to-Site connection in the Azure portal." Step 4, "Create a local network gateway," details the requirement for the public IP address of the on-premises VPN device in the 'IP address' field.

Microsoft Documentation Azure VPN Gateway: "About BGP with Azure VPN Gateway." This documentation explains how BGP can be used to propagate routes between Azure and on-premises networks, enabling scenarios like the transitive routing described in the first statement.

Question: 35

HOTSPOT You have an Azure subscription that contains multiple virtual machine scale sets and multiple Azure load balancers. The load balancers balance traffic across the scale sets. You plan to deploy Azure Front Door to load balance traffic across the load balancers. You need to identify which Front Door SKU to configure, and what to use to route the traffic to the load balancers. The solution must minimize costs. What should you identify? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Answer:

SKU: Premium

Use: Azure Private Link

Explanation:

The goal is to use Azure Front Door to load balance traffic to backend Azure load balancers. The key to this question lies in evaluating the provided options for routing traffic.

• Routing Method: Of the options provided, Azure Private Link is the modern, secure method for connecting Azure Front Door to backend services (like an Azure Load Balancer) without exposing them to the public internet. It ensures traffic remains on the Microsoft private backbone network. The other options, Azure Route Server and service endpoints, are not used for this purpose.

• SKU Requirement: The ability to connect to an origin using Azure Private Link is an exclusive feature of the Premium SKU. The Standard SKU does not support this functionality.

While the prompt includes a requirement to minimize costs, the provided routing options force the selection of a specific architecture. Since Azure Private Link is the only viable and secure routing mechanism listed, you must select the SKU that supports it, which is Premium. Therefore,

Premium becomes the most cost-effective choice that satisfies the implied architectural requirements of the available answers.

References:

Microsoft Azure Documentation - Azure Front Door Tiers: This official document provides a comparison of the Standard and Premium SKUs. The feature table explicitly states that "Private Link support" is available only in the Premium tier.
Reference: Microsoft Docs, "Tier comparison - Azure Front Door," Section: "Feature comparison across tiers."
Microsoft Azure Documentation - Secure origin with Private Link: This guide details how to configure the connection. It states, "You can enable Private Link for your origins such as Application Gateway, virtual machine, internal Load Balancer, and Storage in Azure Front Door Premium tier to lock down your origin from public access."
Reference: Microsoft Docs, "Secure your Origin with Private Link in Azure Front Door," Section: "Overview."