DUMPSARENA

Designing and Implementing Microsoft Azure Networking Solutions

Microsoft AZ-700

Version Demo

Total Demo Questions: 42

Total Premium Questions: 420

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Design, implement, and manage hybrid networking	86
Topic 2, Design and implement core networking infrastructure	146
Topic 3, Design and implement routing	34
Topic 4, Secure and monitor networks	128
Topic 5, New Topic: Contoso Case Study	10
Topic 6, New Topic: Proseware. Inc Case Study	10
Topic 7, Case Study 1	2
Topic 8, Case Study 2	2
Topic 9, Case Study 3	2
Total	420

QUESTION NO: 1

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:

3 is the minimum because a virtual network service endpoint is enabled per subnet and per Azure service type (resource provider), not per individual resource instance. For Azure Storage, you enable the *Microsoft.Storage* service endpoint on the subnet(s) where the virtual machines reside; that single endpoint then applies to any Storage accounts you later secure to that VNet/subnet (such as storage1 and storage2) by configuring their network rules to allow the selected virtual network/subnet. For the database, you enable the service endpoint for the relevant database service type (for example, *Microsoft.Sql* for Azure SQL Database/SQL Server, or the appropriate provider for the DB service shown as DB1) on the same subnet(s). In addition, because service endpoints are configured at the subnet level, if the virtual machines are spread across multiple subnets (as implied by the resource table), you must enable the required service endpoints on each of those subnets. In this scenario, the minimum total count comes from enabling the needed service endpoint types across the necessary subnets, resulting in three service endpoint configurations overall. See: Virtual network service endpoints and Azure Storage network security.

QUESTION NO: 2

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

Name	Type	Description
App1	Azure App Service	A web app
Gateway1	Azure Application Gateway	includes an SSL certificate that has a subject name of *.contoso.com

Gateway1 provides access to App1 by using a URL of http://app1.contoso.com. You create a new web app named App2.

You need to configure Gateway1 to enable minimize administrative effort. What should you configure on Gateway1?

A. a backend pool and a routing

B. a listener and a routing rule

C. a listener, a backend pool, and a rule

D. a listener and a backend pool

ANSWER: B

Explanation:

To publish a new web app through an existing Azure Application Gateway with minimal administrative effort, you should configure a listener and a routing rule. The listener is required to accept incoming requests for the new hostname (for example, app2.contoso.com) and port/protocol combination, and it binds the frontend IP/port to the host header (multi-site hosting). The routing rule is then used to map traffic received by that listener to the appropriate backend target (the web app), typically via a backend pool and HTTP settings referenced by the rule. In Application Gateway, traffic flow is fundamentally "listener → rule → backend," so adding a new site requires at least a new listener (to match the new host name) and a rule (to direct matched requests). This approach avoids reworking existing configuration and cleanly extends the gateway for the additional app.

References: https://learn.microsoft.com/en-us/azure/application-gateway/configuration-overview, https://learn.microsoft.com/en-us/azure/application-gateway/multiple-site-overview

QUESTION NO: 3

You have an Azure virtual network named Vnet1 that hosts an Azure firewall named FW1 and 150 virtual machines. Vnet1 is linked to a private DNS zone named contoso.com. All the virtual machines have their name registered in the contoso.com zone.

Vnet1 connects to an on-premises datacenter by using ExpressRoute.

You need to ensure that on-premises DNS servers can resolve the names in the contoso.com zone.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. On the on-premises DNS servers, configure forwarders that point to the frontend IP address of FW1.

B. On the on-premises DNS servers, configure forwarders that point to the Azure provided DNS service at 168.63.129.16.

C. Modify the DNS server settings of Vnet1.

D. For FW1, enable DNS proxy.

E. For FW1, configure a custom DNS server.

ANSWER: A D

Explanation:

To let on-premises DNS servers resolve records that exist only in an Azure Private DNS zone linked to a virtual network, you need a DNS forwarding path into Azure that can query Azure's internal resolver on behalf of on-premises clients. A common and supported pattern is to use Azure Firewall as a DNS proxy: you enable DNS proxy on the firewall and then configure your on-premises DNS servers to forward the relevant zone (or all unknown queries) to the firewall's private IP address. With DNS proxy enabled, Azure Firewall can accept DNS queries from on-premises over ExpressRoute and forward them to the Azure-provided DNS service for resolution, which includes Private DNS zones linked to the virtual network. This avoids exposing 168.63.129.16 directly from on-premises and centralizes DNS forwarding through a controlled point in Azure. In this scenario, pointing the on-premises forwarders to the frontend IP address of the firewall and enabling DNS proxy on the firewall together provide the required end-to-end name resolution for contoso.com from on-premises.

References: Azure Firewall DNS settings and DNS proxy, Azure Private DNS resolution options

QUESTION NO: 4

You have 10 on-premises networks that are connected by using a 3rd party Software Defined Wide Area Network (SD-WAN) solution. You have an Azure subscription that contains five virtual networks. You plan to connect the Azure virtual networks and the on-premises networks by using an Azure Virtual WAN with a single virtual WAN hub.

You need to ensure that the Azure Virtual WAN can act as a node in the 3rd party SD-WAN solution. What should you include in the solution?

A. An Azure Virtual WAN ExpressRoute gateway

B. A Network Virtual Appliance (NVA)

C. A Site to site gateway (VPN gateway)

D. A Point to site gateway (User VPN gateway)

ANSWER: B

Explanation:

A Network Virtual Appliance (NVA) is the right component to include when you need Azure Virtual WAN to participate as a node in a third-party SD-WAN fabric. In practice, most SD-WAN vendors integrate with Azure by deploying their SD-WAN virtual appliance in Azure and then connecting it to the Virtual WAN hub (typically using hub routing and/or VPN/ExpressRoute connectivity depending on the vendor design). The NVA provides the vendor-specific SD-WAN control/data-plane functions (overlay tunnels, routing policies, segmentation, and orchestration) that Azure Virtual WAN does not natively implement for third-party SD-WAN membership. Azure Virtual WAN provides managed connectivity building blocks (VPN/ER gateways, routing, and hub-to-spoke connectivity), but acting as a "node" in an external SD-WAN solution generally requires the vendor's appliance running in Azure to terminate and participate in the SD-WAN overlay. This is a common pattern described in Azure documentation for integrating NVAs and third-party network virtual appliances with Virtual WAN and hub routing.

References: Microsoft Docs: About Azure Virtual WAN, Microsoft Docs: Virtual WAN routing and routing policies (hub routing/NVA scenarios)

QUESTION NO: 5

You have an Azure virtual network named Vnet1.

You need to ensure that the virtual machines in Vnet1 can access only the Azure SQL resources in the East US Azure region. The virtual machines must be prevented from accessing any Azure Storage resources.

Which two outbound network security group (NSG) rules should you create? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. an allow rule that has the IP address range of Vnet1 as the source and destination of Sq1.EastUS

B. a deny rule that has a source of VirtualNetwork and a destination of Sq1

C. a deny rule that has a source of VirtualNetwork and a destination of 168.63.129.0

D. a deny rule that has the IP address range of Vnet1 as the source and destination of Storage

ANSWER: A D

Explanation:

To meet the requirement, you use NSG outbound rules with Azure service tags so traffic is scoped to specific Azure platform services and (for SQL) to a specific region. Allowing outbound access to Azure SQL in East US is done by creating an outbound allow rule where the source is the virtual network (or the Vnet1 address space) and the destination is the regional SQL service tag (Sql.EastUS). This ensures VMs can reach only Azure SQL endpoints in that region when combined with restrictive deny rules.

To prevent any access to Azure Storage, you add an outbound deny rule from the virtual network (or Vnet1 address space) to the Storage service tag. Service tags represent the published IP ranges for the service and are kept up to date by Microsoft, which is the recommended approach versus manually maintaining IP ranges.

In practice, you would also ensure rule priorities are set so the allow to Sql.EastUS is evaluated before broader denies, and that a final “deny all” (or reliance on the default outbound allow/deny posture as appropriate) results in only the intended destinations being reachable.

References: https://learn.microsoft.com/en-us/azure/virtual-network/service-tags-overview, https://learn.microsoft.com/en-us/azure/virtual-network/network-security-groups-overview

QUESTION NO: 6

Azure virtual networks in the East US Azure region as shown in the following table.

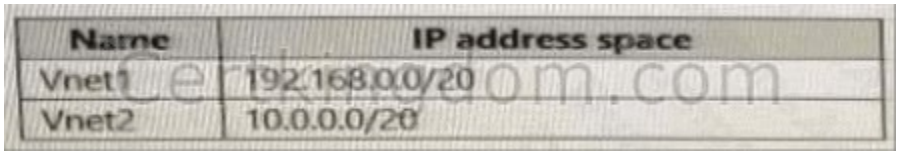

Name	IP address space
Vnet1	192.168.0.0/20
Vnet2	10.0.0.0/20

The virtual networks are peered to one another. Each virtual network contains four subnets.

You plan to deploy a virtual machine named VM1 that will inspect and route traffic between all the

subnets on both the virtual networks.

What is the minimum number of IP addresses that you must assign to VM1?

A. 1

B. 2

C. 4

D. 8

ANSWER: B

Explanation:

2 is the minimum because VM1 must act as a network virtual appliance (NVA) that forwards traffic between two different virtual networks. In Azure, a VM can only be directly connected to a subnet via a network interface (NIC), and each NIC is placed in exactly one subnet within one virtual network. To route traffic between two peered VNets using an NVA, the NVA needs a presence in each VNet so that user-defined routes (UDRs) in each VNet can point to the NVA as the next hop. Practically, that means attaching at least one NIC in a subnet of the first VNet and at least one NIC in a subnet of the second VNet. Each NIC requires at least one private IP configuration, so the minimum number of IP addresses assigned to VM1 is two (one per NIC). You do not need an IP per subnet; instead, you steer traffic from all subnets to the NVA using route tables, and the NVA forwards between networks with IP forwarding enabled on the NICs. For details, see Azure VM multiple NIC behavior and NVA routing with UDRs: https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface?tabs=azure-portal and https://learn.microsoft.com/en-us/azure/virtual-network/virtual-networks-udr-overview.

QUESTION NO: 7 - (DRAG DROP)

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

Name	Type	Location
WebApp1	Web app	West US
VNet1	Virtual network	East US

The IP Addresses settings for Vnet1 are configured as shown in the exhibit.

You need to ensure that you can integrate WebApp1 and Vnet1.

Which three actions should you perform in sequence before you can integrate WebApp1 and Vnet1? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

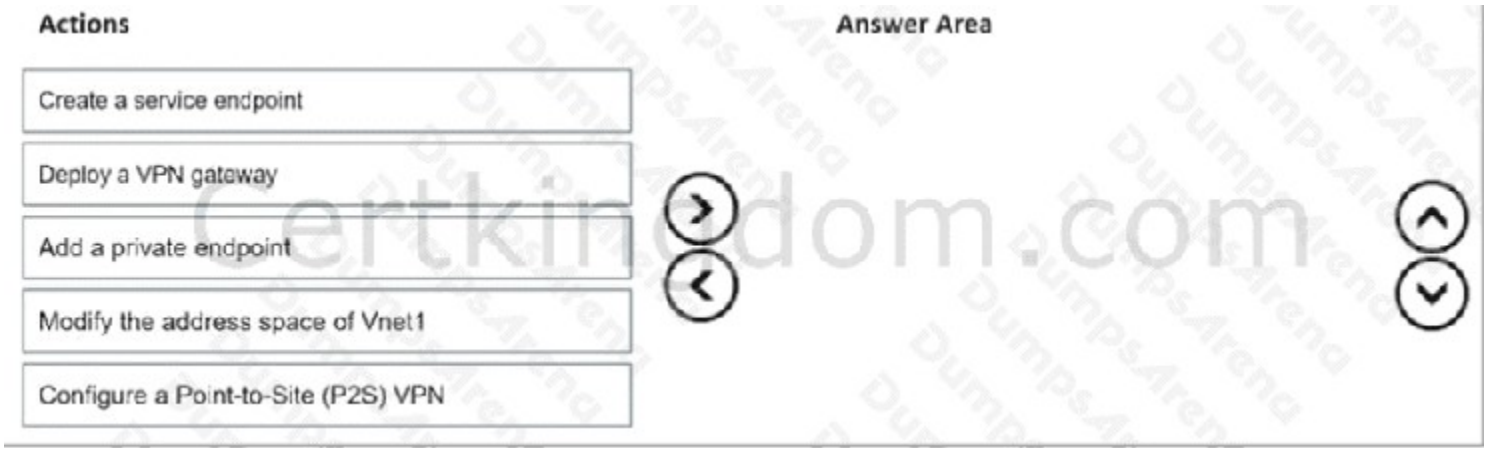

ANSWER:

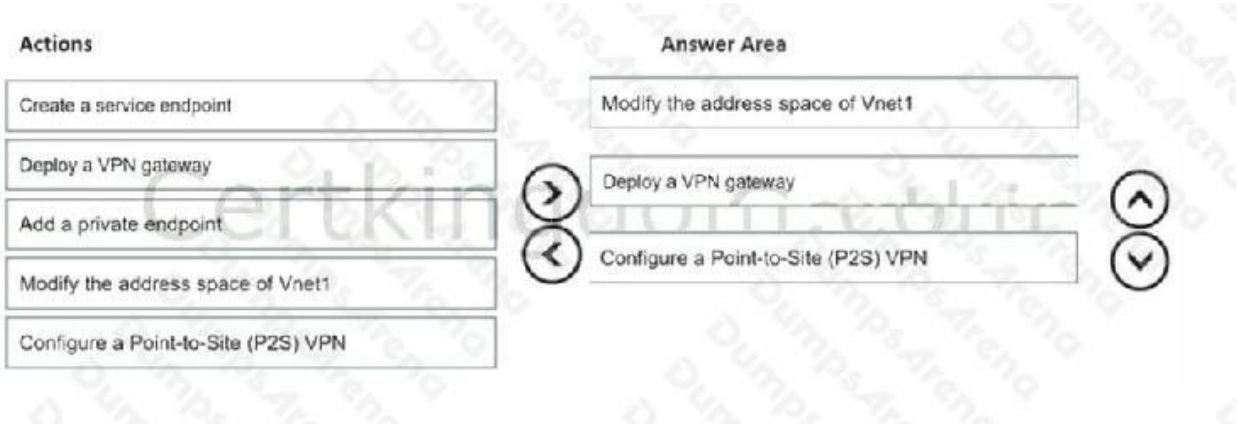

Explanation:

Because WebApp1 is in West US and VNet1 is in East US, you can't use Regional VNet Integration (it requires the app and VNet to be in the same region). In this cross-region scenario, App Service must use the gateway-required VNet integration model, which relies on a Virtual Network Gateway and a Point-to-Site (P2S) configuration so the App Service can establish connectivity into the VNet.

Before you can deploy a Virtual Network Gateway, the VNet must have a dedicated subnet named **GatewaySubnet**. The exhibit shows VNet1 has address space **10.3.0.0/16** and Subnet1 also consumes **10.3.0.0/16**, leaving no free address range to create GatewaySubnet. So the first prerequisite is to modify (expand) the VNet address space so there is room to create GatewaySubnet (for example, by adding another address prefix to the VNet and then creating GatewaySubnet from the newly available space). Once the VNet has room for GatewaySubnet, you can deploy the VPN gateway (Virtual network gateway). After the gateway is deployed, you configure Point-to-Site (P2S) on the gateway, which is the required connectivity method used by gateway-required App Service VNet integration.

This sequence aligns with Microsoft's guidance for gateway-required VNet integration prerequisites and the requirement for a Virtual Network Gateway and P2S setup. See Integrate your app with an Azure virtual network (gateway-required VNet integration) and About VPN Gateway.

QUESTION NO: 8

You are planning the IP addressing for the subnets in Azure virtual networks.

Which type of resource requires IP addresses in the subnets?

A. internal load balancers

B. storage account

C. serviice endpoints

D. service endpoint policies

ANSWER: A

Explanation:

internal load balancers require IP addresses in the subnet because an internal (private) Azure Load Balancer is reached through a frontend IP configuration that uses a private IP from the virtual network. That private frontend IP must be allocated

from the address space of a specific subnet (either dynamically or as a static private IP), which means you must plan subnet capacity to accommodate it. In contrast, resources like storage accounts are platform services and don't consume IPs from your VNet subnets unless you deploy a private endpoint (not mentioned here). Service endpoints and service endpoint policies extend VNet identity/routing to Azure services but do not allocate additional IP addresses in the subnet; they work by updating route/ACL behavior for existing subnet traffic rather than creating new IP-consuming interfaces. When designing subnet IP ranges, you therefore need to account for private frontend IPs used by internal load balancers (and other IP-consuming resources like NICs, private endpoints, etc.).

References: Microsoft Docs: Azure Load Balancer overview, Microsoft Docs: Load Balancer components (frontend IP configuration)

QUESTION NO: 9 - (DRAG DROP)

You have an on-premises network

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains an ExpressRoute gateway named Gateway 1.

You need to implement an ExpressRoute solution from a third-party provider named Fabrikam, Inc. The solution must ensure that devices on the on-premises network can connect to the Azure resources on VNet1.

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.

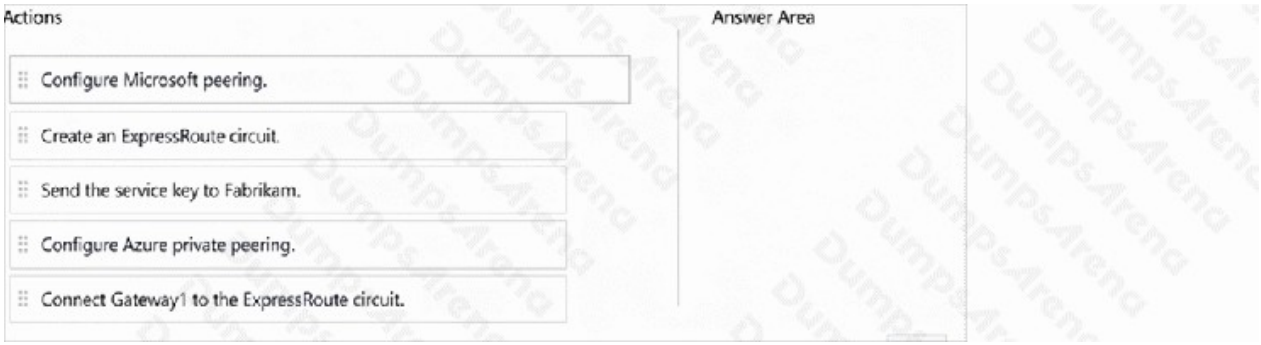

ANSWER:

Explanation:

To connect an on-premises network to Azure resources in VNet1 over ExpressRoute, you follow the typical provider-based provisioning flow. You start in Azure by creating an ExpressRoute circuit. Creating the circuit is what generates the circuit's service key (also called the service key or service identifier), which is the value the connectivity provider needs to associate their physical/virtual cross-connect with your Azure circuit.

Next, you send that service key to Fabrikam so they can provision the ExpressRoute connectivity on their side. Until the provider completes provisioning, the circuit won't reach the expected state for end-to-end connectivity.

After the provider provisions the circuit, you configure *Azure private peering* on the ExpressRoute circuit. Private peering is the peering type that enables private IP connectivity between your on-premises network and Azure virtual networks (including VNet1). This step includes supplying the BGP peer IPs, ASN, VLAN ID (if required), and related settings so routes can be exchanged privately.

Finally, you connect the existing ExpressRoute gateway (Gateway1) in VNet1 to the ExpressRoute circuit by creating the circuit-to-gateway connection. This is what actually links the VNet to the circuit so that learned routes can propagate and traffic can flow between on-premises and VNet1 resources.

Microsoft peering isn't required for this goal because it's intended for reaching Microsoft public services, not for private VNet connectivity. For more details, see Create and modify an ExpressRoute circuit and Configure routing (peerings) for an ExpressRoute circuit.

QUESTION NO: 10 - (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

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.

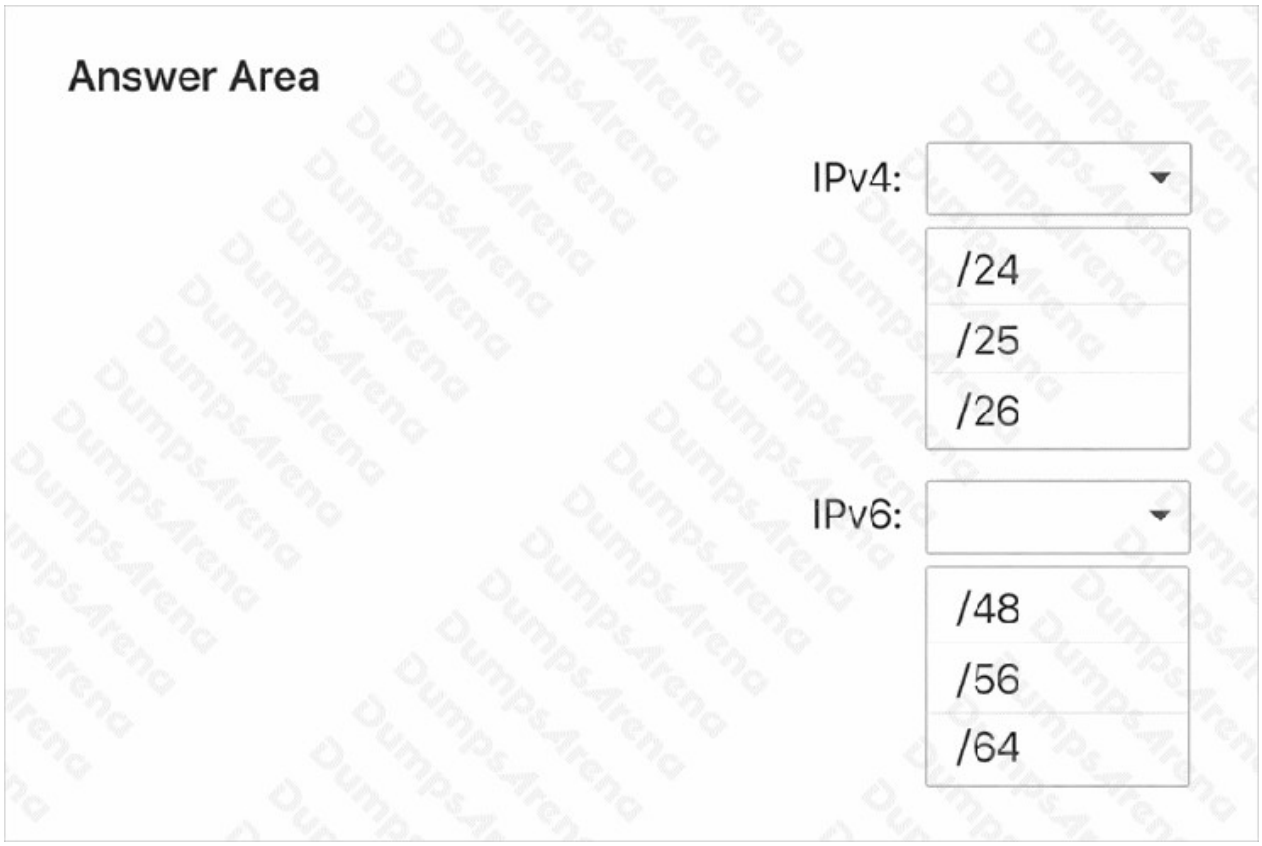

ANSWER:

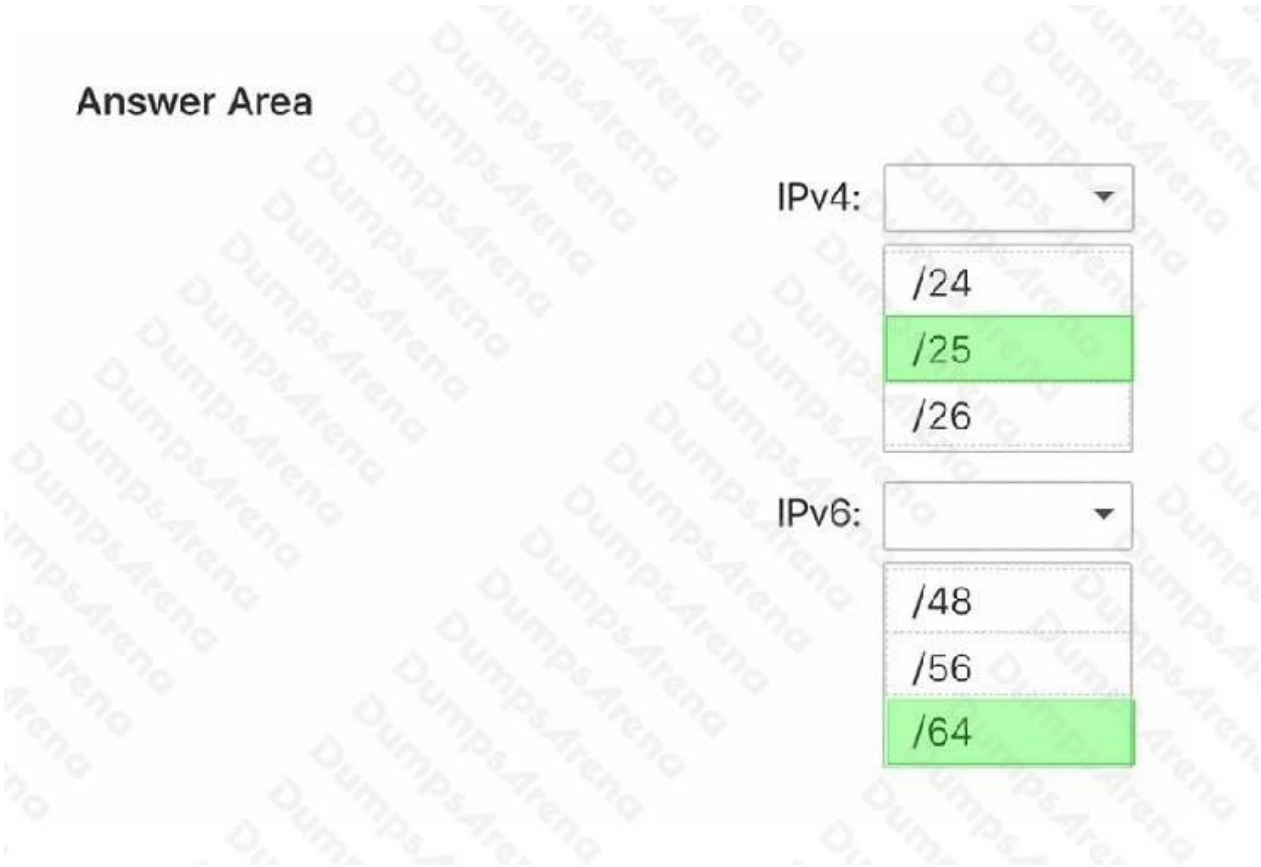

Explanation:

To meet the requirements, you need an address space that's large enough to carve out a dedicated **GatewaySubnet** for the VPN gateway and still leave room for one or more VM subnets, while also keeping the overall allocation as small as possible to "maximize usable IP addresses" (in practice, that means choosing the smallest prefix that still supports the design).

For **IPv4**, Azure VPN Gateway requires a **GatewaySubnet** that is at least **/27** (and Microsoft commonly recommends /27 or larger for growth). If you chose an IPv4 VNet address space of **/26**, you would not be able to create a /27 GatewaySubnet inside it (a subnet can't be larger than the VNet's address space), and you also need additional space for VM subnets. The smallest option that still allows you to create a /27 (or larger) GatewaySubnet and additional VM subnets is **/25**. That satisfies the VPN gateway requirement and preserves the most efficiency compared to /24.

For **IPv6**, Azure virtual network IPv6 subnets are expected to be **/64** in typical deployments, and /64 is the standard subnet size used for IPv6 addressing in Azure VNets. Selecting **/64** also aligns with the goal of using the most specific (smallest) prefix available in the dropdown while still supporting resource deployment.

References: About VPN Gateway and IPv6 for Azure Virtual Network.

QUESTION NO: 11 - (SIMULATION)

Task 2

You need to create an Azure Firewall instance named FW1 that meets the following requirements: Has an IP address from the address range of 10.1.255.0

Uses a new Premium firewall policy named FW-pohcy1 Routes traffic directly to the internet

ANSWER: See the explanation for the answer

Explanation:

instructions.

To create an Azure Firewall instance, you need to go to the Azure portal and select Create a resource. Type firewall in the search box and press Enter. Select Firewall and then select Create1.

To assign an IP address from the address range of 10.1.255.0 to the firewall, you need to select a public IP address that belongs to that range. You can either create a new public IP address or use an existing one1.

To use a new Premium firewall policy named FW-policy1, you need to select Premium as the Firewall tier and create a new policy with the name FW-policy12. A Premium firewall policy allows you to configure advanced features such as TLS Inspection, IDPS, URL Filtering, and Web Categories3.

To route traffic directly to the internet, you need to enable SNAT (Source Network Address Translation) for the firewall. SNAT allows the firewall to use its public IP address as the source address for outbound traffic4.

QUESTION NO: 12

You have an Azure virtual network and an on-premises datacenter.

You are planning a Site-to-Site VPN connection between the datacenter and the virtual network.

Which two resources should you include in your plan? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a user-defined route

B. a virtual network gateway

C. Azure Firewall

D. Azure Web Application Firewall (WAF)

E. an on-premises data gateway

F. an Azure application gateway

G. a local network gateway

ANSWER: B G

Explanation:

For an Azure Site-to-Site (S2S) VPN, you must plan for the two gateway resources that represent each side of the VPN connection in Azure. A virtual network gateway is the Azure-side VPN endpoint that gets deployed into a dedicated GatewaySubnet in the virtual network and provides the IPsec/IKE termination for the tunnel. A local network gateway is the Azure resource that represents your on-premises VPN device and network: it stores the on-premises public IP address of the VPN device and the on-premises address prefixes that should be reachable over the tunnel. With these two resources in place, you can then create the VPN connection object between them and configure matching IPsec/IKE settings on the on-

premises device. These are the core required resources for an S2S VPN design; other items like firewalls, WAF, application gateways, or user-defined routes may be used in specific architectures but are not required to establish the S2S VPN itself.

References: Tutorial: Create a site-to-site VPN connection in the Azure portal, About VPN Gateway

QUESTION NO: 13

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:

Increase the number of NAT IP addresses assigned to PL1. A Private Link service uses a Standard Load Balancer, and outbound connections from the service provider side can be constrained by source NAT (SNAT) port availability. Each outbound flow consumes SNAT ports from the NAT IP(s) used for outbound translation; as concurrent outbound connections grow, you can exhaust available ephemeral ports and see outbound connection failures or throttling. By adding more NAT IP addresses to the Private Link service, you effectively increase the pool of available SNAT ports, allowing PL1 to scale outbound connection capacity and support a higher volume of outbound traffic. This approach aligns with Azure's guidance for scaling outbound connectivity by increasing the number of outbound public IPs (or NAT IPs) used for SNAT, which increases the number of available ports and reduces the risk of SNAT exhaustion. For details on Private Link service NAT IP configuration and scaling, see Microsoft Docs: Private Link service and for SNAT/port exhaustion concepts and scaling outbound, see Microsoft Docs: Load Balancer outbound connections.

QUESTION NO: 14

You have an Azure application gateway named AGW1 that has a routing rule named Rule1. Rule 1 directs traffic for http://www.contoso.com to a backend pool named Pool1. Pool1 targets an Azure virtual machine scale set named VMSS1.

You deploy another virtual machine scale set named VMSS2.

You need to configure AGW1 to direct all traffic for http://www.adatum.com to VMSS2.

The solution must ensure that requests to http://www.contoso.com continue to be directed to Pool1.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Add a backend pool.

B. Modify an HTTP setting.

C. Add an HTTP setting.

D. Add a listener.

E. Add a rule.

ANSWER: A D E

Explanation:

To route a second hostname (http://www.adatum.com) to a different backend while keeping the existing hostname (http://www.contoso.com) working as-is, Application Gateway needs a separate set of configuration objects for that new traffic flow. First, you add a backend pool that contains VMSS2 (or its IPs/targets), because backend pools define where the gateway can send traffic. Next, you add a listener configured for the new host name (www.adatum.com). In Application Gateway, the listener is what accepts incoming requests and can match on host header (multi-site hosting), so a distinct listener is required to differentiate adatum.com from contoso.com. Finally, you add a rule that binds the new listener to the new backend pool (and an HTTP setting). This new rule ensures traffic for www.adatum.com is routed to VMSS2, while the existing Rule1 continues to route www.contoso.com to Pool1 without being impacted. This aligns with the standard Application Gateway request-routing model of listener + rule + backend pool (+ HTTP settings). See Application Gateway configuration overview and Multiple site hosting with Application Gateway.

QUESTION NO: 15

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 the following resources:

- A virtual network named Vnet1
- A subnet named Subnet1 in Vnet1
- A virtual machine named VM1 that connects to Subnet1
- Three storage accounts named storage1, storage2, and storage3

You need to ensure that VM1 can access storage1. VM1 must be prevented from accessing any other storage accounts.

Solution: You create a network security group (NSG) and associate the NSG to Subnet1.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes, using a network security group associated to Subnet1 can meet the goal because NSGs can control outbound traffic from VM1 at the subnet (or NIC) level. Azure Storage is reached over public endpoints (unless you use private endpoints), and NSG outbound rules can be written to allow traffic only to the specific destination used by storage1 and deny traffic to other storage accounts. In practice, this is typically done by adding an outbound allow rule for the Azure Storage service tag scoped to the region and then further restricting access using storage account firewall settings (selected networks) so only the intended path is permitted. The key point is that NSGs are the Azure construct designed to allow/deny traffic flows, and when applied to the subnet they affect all VMs in that subnet, including VM1, enabling you to enforce “allow to storage1, block to others” at the network layer.

For details on how NSG rules work and how service tags like Storage can be used in NSG rules, see Network security groups overview and Virtual network service tags.

QUESTION NO: 16

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

Name	Description
App1	Web app
FD1	Azure Front Door Premium profile that has an endpoint and an origin group

You need to configure a solution to meet the following requirements; App1 must be assigned a private endpoint

Access to App1 from the internet must be routed via FD1. What should you configure on FD1?

A. a rule that has the route configuration override action

B. a route that redirects traffic

C. an origin that enables the Azure Private Link service

D. a security policy that redirects traffic

ANSWER: C

Explanation:

To keep App1 private (reachable only through a private endpoint) while still allowing internet users to access it through FD1, you must configure Azure Front Door to connect to the backend using Private Link. In Azure Front Door Premium, this is done at the origin level by enabling the Azure Private Link service for the origin that represents App1. When Private Link is enabled on the origin, Front Door establishes private connectivity from the Front Door edge to the backend via an approved Private Link connection, so the app can remain non-public while Front Door continues to serve as the internet-facing entry point. This design meets both requirements: App1 is assigned a private endpoint, and all internet access is routed through FD1 rather than directly to the app’s public endpoint. This is the supported pattern for exposing private origins through Front Door Premium using Private Link integration.

References: Azure Front Door Premium with Private Link, Configure origins and origin groups in Azure Front Door

QUESTION NO: 17

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

Name	Type	Description
VNet1	Virtual network	Contains a subnet named Subnet1
storage1	Storage account	*None*
VM1	Virtual machine	Linked to Subnet1
VM2	Virtual machine	Linked to Subnet1

You need to ensure that VM1 and VM2 can connect only to storage1. The solution must meet the following requirements:

Prevent VM1 and VM2 from accessing any other storage accounts. Ensure that storage1 is accessible from the internet.

What should you use?

A. a network security group (NSG)

B. a private endpoint

C. a private link

D. a service endpoint policy

ANSWER: D

Explanation:

Using a service endpoint policy is the right fit because it lets you restrict Azure Storage access over a virtual network service endpoint to only specific storage accounts. In this scenario, you enable the Microsoft.Storage service endpoint on the subnet that hosts VM1 and VM2, then apply a service endpoint policy that explicitly allows only storage1. This prevents the VMs from reaching any other storage accounts via the service endpoint path, meeting the requirement to block access to other storage accounts while still allowing access to storage1.

At the same time, storage1 can remain accessible from the internet because service endpoint policies don't require you to disable public network access on the storage account. They control what the subnet can reach over the service endpoint, not whether the storage account has a public endpoint. This combination satisfies both requirements: tightly scoped storage access from the VMs and continued internet accessibility for storage1.

References: https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-service-endpoint-policies-overview, https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-service-endpoints-overview

QUESTION NO: 18

You have an Azure subscription that is linked to an Azure AD tenant named contoso.onmicrosoft.com. The subscription contains the following resources:

• A virtual network named Vnet1

• An App Service plan named ASPI

• An Azure App Service named webapp1

• An Azure private DNS zone named private.contoso.com

• Virtual machines on Vnet1 that cannot communicate outside the virtual network

You need to ensure that the virtual machines on Vnet1 can access webapp1 by using a URL of https:/Avwwprivate.contosocom.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Create a private endpoint for webapp1.

B. Create a service endpoint for webapp1.

C. Create a CNAME record that maps www.pnvate.contoso.com to webapp1.privatelink.azurewebsites.net.

D. Create a CNAME record that maps wwwprivatemntoso.com to webapp1.contoso.onmicrosoft.com.

E. Register an enterprise application in Azure AD for webapp1.

F. Create a CNAME record that maps wow.private.contoso.com to webapp 1 private@ntoso.com.

ANSWER: A C

Explanation:

To allow virtual machines that have no outbound connectivity to reach an Azure App Service privately, you use Azure Private Link for App Service. Creating a private endpoint for the app injects a private IP address into a subnet in Vnet1 and maps the App Service's inbound traffic to that private IP, so the VMs can connect without using the public internet. With Private Link, name resolution must also return the private endpoint IP. The standard approach is to create a DNS record for the desired hostname (www.private.contoso.com) that aliases to the app's privatelink FQDN (webapp1.privatelink.azurewebsites.net). When the private DNS zone is linked to the virtual network, the VMs resolve the hostname to the private endpoint and can access the app using the requested URL. This aligns with Microsoft guidance that App Service private endpoints require private DNS integration and typically use CNAME aliasing to the privatelink zone name for custom hostnames. See Azure App Service private endpoint and Private endpoint DNS configuration.

QUESTION NO: 19

You plan to configure BGP for a Site-to-Site VPN connection between a datacenter and Azure.

Which two Azure resources should you configure? Each correct answer presents a part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

A. a virtual network gateway

B. Azure Application Gateway

C. Azure Firewall

D. a local network gateway

E. Azure Front Door

ANSWER: A D

Explanation:

To use BGP over an Azure Site-to-Site VPN, you configure BGP on the Azure VPN endpoint and define the on-premises BGP peer details. In Azure, the VPN endpoint is provided by a virtual network gateway (VPN Gateway). This gateway must be created with a gateway SKU that supports BGP, and you then enable BGP and set the Azure-side ASN and BGP peering address (typically from the GatewaySubnet). You also need a local network gateway to represent the on-premises VPN device and its address space; when using BGP, the local network gateway is additionally where you specify the on-premises BGP peer IP address and ASN so Azure can establish the BGP session across the IPsec tunnel. Together, these two resources provide the required configuration objects for BGP routing exchange between Azure and your datacenter over the Site-to-Site VPN connection. For implementation details and required properties (ASN, BGP peer address, and enabling BGP on the gateway), see Azure VPN Gateway BGP overview and Configure BGP for Azure VPN Gateway.

QUESTION NO: 20

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 on Azure subscription that contains on Azure Virtual WAN named VWAN1.VWAN1 contains a hub named Hub1.

Hub! has a security status of Unsecured.

You need to ensure that the security status of Hub1 is marked as Secured. Solution: You implement Azure Firewall.

Does this meet the requirement?

A. Yes

B. No

ANSWER: A

Explanation:

Yes meets the requirement because a Virtual WAN hub's "Secured" status is achieved by enabling hub security through Azure Firewall (specifically, Azure Firewall deployed as a secured virtual hub, sometimes referred to as Azure Firewall in a Virtual WAN hub). In Virtual WAN, the hub security provider is Azure Firewall, and when you deploy/enable Azure Firewall in the hub, the hub becomes a secured virtual hub and the portal reflects the hub security status as secured. This is the intended mechanism for adding centralized, managed network security controls (stateful filtering, application and network rules, threat intelligence, etc.) directly in the Virtual WAN hub so that traffic traversing the hub can be inspected and controlled. Therefore, implementing Azure Firewall in the hub is the correct action to change the hub security status from Unsecured to Secured.

References: Configure Azure Firewall in an Azure Virtual WAN hub, What is Azure Virtual WAN?

QUESTION NO: 21

You plan to deploy an Azure virtual network.

You need to design the subnets.

Which three types of resources require a dedicated subnet? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. VPN gateway

B. Azure Bastion

C. Azure Active Directory Domain Services (Azure AD DS)

D. Azure Application Gateway v2

E. Azure Private Link

ANSWER: A B D

Explanation:

Several Azure networking services must be deployed into their own dedicated subnet because the platform requires specific subnet names, reserved IP usage, and service-managed scaling behavior that can't safely coexist with general workloads. A VPN gateway is deployed into a subnet named **GatewaySubnet**, which must be dedicated to the gateway so Azure can allocate addresses and manage gateway instances correctly. Azure Bastion similarly requires its own dedicated subnet named **AzureBastionSubnet** to host the Bastion service and its scaling infrastructure. Azure Application Gateway v2 must also be placed in a dedicated subnet; this is a documented requirement to ensure the gateway's managed components, routing, and autoscaling operate without conflicts from other resources in the same subnet.

These dedicated-subnet requirements are core design considerations when planning address space and subnetting for an Azure virtual network, especially to avoid later rework (for example, needing to resize or split subnets). For official guidance, see Microsoft's documentation on integrating Azure services with virtual networks and the specific service deployment requirements for Bastion, VPN Gateway, and Application Gateway.

References: https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-for-azure-services, https://learn.microsoft.com/en-us/azure/bastion/bastion-overview

QUESTION NO: 22

You have an Azure subscription that contains the public IP addresses shown in the following table.

Name	IP version	SKU	IP address assignment
IP1	IPv4	Basic	Static
IP2	IPv4	Basic	Dynamic
IP3	IPv4	Standard	Static
IP4	IPv6	Basic	Dynamic
IP5	IPv6	Standard	Static

You plan to deploy a NAT gateway named NAT1.

Which public IP addresses can be used as the public IP address for NAT1?

A. IP3 only

B. IP5 only

C. IP2 and IP4 only

D. IP1, IP3 and IP5 only

E. IP3 and IP5 only

ANSWER: A

Explanation:

An Azure NAT gateway can only be associated with public IP resources that are IPv4, use the Standard SKU, and are configured as Static allocation. This is because NAT gateway is a zonal, highly available service that relies on the capabilities of Standard public IPs (including zone support and predictable behavior). Basic SKU public IP addresses aren't supported for NAT gateway, and IPv6 public IP addresses can't be used because NAT gateway provides outbound SNAT for IPv4 traffic only. Therefore, from the listed public IPs, the only eligible address is the one that matches all three requirements: Standard SKU + Static + IPv4. In the table, that corresponds to IP3, so IP3 is the only public IP address that can be used for NAT1.

References: Azure NAT Gateway overview, NAT gateway resource and requirements

QUESTION NO: 23

You have an Azure virtual network named Vnet1 that hosts an Azure firewall named FW1 and 150 virtual machines. Vnet1 is linked to a private DNS zone named contoso.com. All the virtual machines have their name registered in the contoso.com zone.

Vnet1 connects to an on-premises datacenter by using ExpressRoute.

You need to ensure that on-premises DNS servers can resolve the names in the contoso.com zone. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. On the on-premises DNS servers, configure forwarders that point to the frontend IP address of FW1.

B. On the on-premises DNS servers, configure forwarders that point to the Azure provided DNS service at 168.63.129.16.

C. Modify the DNS server settings of Vnet1.

D. For FW1, enable DNS proxy.

E. For FW1, configure a custom DNS server.

ANSWER: A D

Explanation:

To allow on-premises DNS servers to resolve records that exist only in an Azure Private DNS zone linked to a virtual network, you need a DNS forwarding path into Azure that can query Azure's internal resolver for that virtual network. A common and supported pattern is to use Azure Firewall as a DNS forwarder by enabling DNS proxy on the firewall and then configuring the on-premises DNS servers to forward the relevant zone (or all unknown queries) to the firewall's private IP address. With DNS proxy enabled, Azure Firewall can accept DNS queries from on-premises and forward them to the Azure-provided DNS service, which is able to resolve names in the linked private DNS zone. This avoids exposing the Azure resolver directly to on-premises and provides a centralized, controllable forwarding point in Azure. In practice, you configure conditional forwarders on-premises for contoso.com to the firewall IP, and ensure DNS proxy is enabled so the firewall will relay those queries appropriately. This is the documented approach for integrating on-premises workloads with Azure Private DNS zones using a DNS forwarder and aligns with Azure Firewall's DNS proxy capability.

References: https://learn.microsoft.com/en-us/azure/private-link/private-endpoint-dns#on-premises-workloads-using-a-dns-forwarder, https://learn.microsoft.com/en-us/azure/firewall/dns-settings

QUESTION NO: 24

You have an Azure subscription that contains an Azure Front Door named FD1. FD1 is configured as shown in the following exhibit.

You need to enable Azure Private Link for FD1. What should you do first?

A. Create an origin group.

B. Add an endpoint.

C. Change Pricing Tier to Azure Front Door Premium.

D. Create a custom route.

ANSWER: C

Explanation:

To enable Azure Private Link with Azure Front Door, you must be using Azure Front Door Standard or Premium (Azure Front Door (classic) doesn't support Private Link, and Private Link for origins is a Premium capability). Therefore, the first required step is to move the profile to a tier that supports Private Link—specifically Azure Front Door Premium—so that you can configure Private Link on an origin (for example, to reach an Azure App Service, Storage, or other supported private endpoint-enabled service) and have Front Door connect privately over Microsoft's backbone. After the tier supports it, you can then configure origin settings to use Private Link and approve the private endpoint connection on the origin resource. Without upgrading the pricing tier, the portal/API won't expose the Private Link configuration options for the origin, so no subsequent configuration (routes, endpoints, origin groups) will achieve the requirement. This is why changing the pricing tier is the necessary first action before any Private Link-specific configuration can be performed.

References: Azure Front Door (Standard/Premium) - Private Link, Azure Front Door tiers and features

QUESTION NO: 25

You have Azure App Service apps in the West US Azure region as shown in the following table.

Name	App Service Plan	Number of instances
App1	ASP1	3
App2	ASP1	3
App3	ASP2	2
App4	ASP3	1

You need to ensure that all the apps can access the resources in a virtual network named Vnet1 without forwarding traffic through the internet.

How many integration subnets should you create?

A. 0

B. 1

C. 3

D. 4

E. 6

ANSWER: B

Explanation:

1

For Azure App Service VNet Integration (the feature used to let web apps reach resources in an Azure virtual network privately), the integration is configured at the app level but it consumes an *integration subnet* in the target virtual network. A key design constraint is that an integration subnet can be used by multiple apps only when those apps are in the same App Service plan. In other words, the subnet is effectively shared per App Service plan: once a plan is integrated with a subnet, any other apps in that same plan can use the same integration subnet to reach VNet resources without sending traffic over the public internet.

Therefore, you should create one integration subnet for each distinct App Service plan represented in the table. Since the table indicates there is only a single App Service plan hosting the apps, only one integration subnet is required to enable all of them to access resources in Vnet1 privately.

References: Azure App Service VNet Integration, Enable VNet Integration

QUESTION NO: 26

You have an Azure application gateway named AGW1 that has a routing rule named Rule1. Rule 1 directs traffic for http://www.contoso.com to a backend pool named Pool1. Pool1 targets an Azure virtual machine scale set named VMSS1.

You deploy another virtual machine scale set named VMSS2.

You need to configure AGW1 to direct all traffic for http://www.adatum.com to VMSS2.

The solution must ensure that requests to http://www.contoso.com continue to be directed to Pool1. Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Add a backend pool.

B. Modify an HTTP setting.

C. Add an HTTP setting.

D. Add a listener.

E. Add a rule.

ANSWER: A D E

Explanation:

To send traffic for a second host name to a different virtual machine scale set while keeping existing host-based routing intact, you add a new set of Application Gateway components for that host. You first create a new backend pool that contains VMSS2 as the target so the gateway has a distinct destination for http://www.adatum.com. Next, you add a new listener configured for the adatum host name (host-based listener), which is how Application Gateway differentiates requests for http://www.adatum.com from those for http://www.contoso.com. Finally, you add a new request-routing rule that links the new listener to the new backend pool (and associated HTTP settings). This preserves the existing Rule1 mapping for contoso to Pool1/VMSS1 while introducing a separate mapping for adatum to VMSS2. These are the core building blocks of

host-based routing on Application Gateway: listener + rule + backend pool (plus HTTP settings as needed). See Application Gateway configuration overview and multiple-site hosting (host-based routing).

QUESTION NO: 27

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 Front Door Premium profile named AFD1 and an Azure Web Application Firewall (WAF) policy named WAF1. AFD1 is associated with WAFT.

You need to configure a rate limit for incoming requests to AFD1. Solution: You configure a custom rule for WAF1.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes, configuring a custom rule in the Azure Web Application Firewall policy meets the goal because rate limiting for Azure Front Door is implemented through WAF custom rules. In Front Door Premium (and Standard), WAF policies support custom rules that can use the *Rate limit* rule type to count requests that match specified conditions (for example, matching all requests, a specific path, host, or client IP) and then take an action such as Block when the threshold is exceeded within a defined duration. Since AFD1 is associated with the WAF policy, the rate-limiting behavior is enforced at the Front Door edge for incoming requests that match the rule, which is exactly what's required to limit request rates to AFD1. This approach is the documented method for applying rate limiting in Azure Front Door WAF and is managed centrally in the WAF policy rather than requiring changes to the backend application.

References: Azure Front Door WAF custom rules, Azure Front Door WAF rate limiting

QUESTION NO: 28

You have an Azure application gateway named AGW1 that has a routing rule named Rule1. Rule 1 directs traffic for http://www.contoso.com to a backend pool named Pool1. Pool1 targets an Azure virtual machine scale set named VMSS1.

You deploy another virtual machine scale set named VMSS2.

You need to configure AGW1 to direct all traffic for http://www.adatum.com to VMSS2.

The solution must ensure that requests to http://www.contoso.com continue to be directed to Pool1.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Add a backend pool.

B. Modify an HTTP setting.

C. Add an HTTP setting.

D. Add a listener.

E. Add a rule.

ANSWER: A D E

Explanation:

To route a second hostname to a different backend while keeping the existing hostname behavior intact, Application Gateway needs a separate set of routing components for the new site. First, you add a backend pool that targets VMSS2 so the gateway has a distinct backend destination for requests to http://www.adatum.com. Next, you add a listener configured for the new host name (www.adatum.com). Hostname-based routing in Application Gateway is implemented by associating a listener with a specific host name (and frontend IP/port), which allows the gateway to distinguish traffic for www.adatum.com from www.contoso.com even if both use the same frontend port. Finally, you add a rule that links the new listener to the new backend pool (and an HTTP setting). This rule ensures requests matching the www.adatum.com listener are forwarded to VMSS2, while the existing Rule1 continues to send www.contoso.com traffic to Pool1/VMSS1. This is the standard pattern for multi-site hosting on Application Gateway: listener per site/host name plus a corresponding routing rule to the appropriate backend pool.

References: Application Gateway configuration overview, Multiple site hosting on Application Gateway

QUESTION NO: 29

You have an Azure subscription that contains the Azure app service web apps show in the following table:

Name	Location	Description
App1eu	West Europe	Production app service for a URL of https://www.fabrikam.com
App1us	East US	Standby app service for a URL of https://www.fabrikam.com

You need to deploy Azure Traffic Manager. The solution must meet the following requirements: Traffic to https//www.fabrikam.com must be directed to App1eu.

If App1eu becomes unresponsive, all the traffic to https://www.fabrikam.com must be directed to App1us. You need to implement Traffic Manager to meet the requirements.

Which two resources should you create? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. a Traffic Manager profile that uses the priority routing method

B. a Traffic Manager profile that uses the geographic routing method

C. a CNAME record in a DNS domain named fabrikam.com

D. a TXT record in a DNS domain named tabrikam.com

E. a real user measurements key in Traffic Manager

ANSWER: A C

Explanation:

To ensure requests for https://www.fabrikam.com are sent to a preferred endpoint and automatically fail over to a secondary endpoint when the preferred one becomes unhealthy, you should use Azure Traffic Manager with the priority routing method. Priority routing is designed specifically for active/passive scenarios: Traffic Manager continuously probes the primary endpoint (App1eu) using its health checks, and if the endpoint becomes unresponsive or is marked degraded, Traffic Manager directs all new DNS responses to the next-priority endpoint (App1us). This meets the requirement for deterministic primary routing with automatic failover.

Because Traffic Manager is DNS-based, you must also map your custom host name (www.fabrikam.com) to the Traffic Manager profile's DNS name. The standard way to do this is to create a CNAME record in the fabrikam.com DNS zone that points www to the Traffic Manager profile (for example, <profile>.trafficmanager.net). This allows clients resolving www.fabrikam.com to receive Traffic Manager's DNS answers and be routed according to priority and health. See Traffic Manager routing methods and Create and manage a Traffic Manager profile.

QUESTION NO: 30

Your on-premises network contains a DNS server named Server 1.

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

Name	Type	Description
VNet1	Virtual network	*None*
VM1	Virtual machine	Connected to VNet1 Connected to storage1 by using a private endpoint
storage1	Storage account	*None*

The on-premises network is connected to VNet1 by using a Site-to-Site (S2S) VPN.

You need to ensure that Server1 can resolve the DNS name of storage1. The solution must minimize costs and administrative effort.

What should you use?

A. an Azure Private DNS zone

B. an Azure virtual machine that hosts a DNS service

C. an Azure public DNS zone

D. Azure DNS Private Resolver

ANSWER: D

Explanation:

Azure DNS Private Resolver is the right choice because it provides a managed way to enable DNS resolution between on-premises networks and Azure private DNS zones over private connectivity such as a Site-to-Site VPN. In this scenario, the

storage account name (storage1) is expected to resolve to a private endpoint IP in VNet1, which is typically published in an Azure Private DNS zone (for example, privatelink.blob.core.windows.net). On-premises DNS servers don't automatically have visibility into Azure private DNS zones, so you need a forwarding path that can answer those queries from Azure.

With Azure DNS Private Resolver, you deploy inbound endpoints in VNet1 so that Server1 can forward queries for the relevant private DNS zones to Azure, and Azure will resolve them using the linked private DNS zone records. This avoids deploying and maintaining your own DNS VM infrastructure and reduces administrative overhead while keeping name resolution private and consistent across hybrid networks.

References: Azure DNS Private Resolver overview, Private endpoint DNS configuration

QUESTION NO: 31

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 application gateway that has Azure Web Application Firewall (WAF) enabled.

You configure the application gateway to direct traffic to the URL of the application gateway.

You attempt to access the URL and receive an HTTP 403 error. You view the diagnostics log and discover the following error.


```
{
 "timeStamp": "2021-06-02T18:13:45+00:00",
 "resourceID": "/SUBSCRIPTIONS/489f2hht-se7y-907v-g571-463hw3679512/RESOURCEGROUPS/RG1/PROVIDERS/MICROSOFT.NETWORK/APPLICATIONGATEWAYS/AGW1",
 "operationName": "ApplicationGatewayFirewall",
 "category": "ApplicationGatewayFirewallLog",
 "properties": {
   "instanceId": "appgw_0",
   "clientIp": "137.135.10.24",
   "clientPort": "",
   "requestUri": "/login",
   "ruleSetType": "OWASP_CRS",
   "ruleSetVersion": "3.0.0",
   "ruleId": "920300",
   "message": "Request Missing an Accept Header",
   "action": "Matched",
   "site": "Global",
   "details": {
     "message": "Warning. Match of \\\"pm AppleWebKit Android\\\" against \\\"REQUEST_HEADER:User-Agent\\\" required. ",
     "data": "",
     "file": "rules\/REQUEST-920-PROTOCOL-ENFORCEMENT.conf",
     "line": "1247"
   },
   "hostname": "app1.contoso.com",
   "transactionId": "f7546159y1hjk7wa114568if5131t68h7",
   "policyId": "default",
   "policyScope": "Global",
   "poplicyScopeName": "Global",
 }
}
```


You need to ensure that the URL is accessible through the application gateway.

Solution: You create a WAF policy exclusion for request headers that contain 137.135.10.24.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

“No” is correct because creating a WAF policy exclusion for request headers that contain 137.135.10.24 is not an appropriate or reliable way to make the application gateway URL accessible. WAF exclusions are designed to skip inspection for specific request elements (such as a particular header name, cookie name, query string argument, or request body field) when those elements are known to cause false positives in managed rule evaluation. They are not intended to “allowlist” a client IP address, nor do they generally match against arbitrary header values unless that header is the one being evaluated by a rule and the exclusion is scoped correctly to the rule set/rule group. If access is being blocked with HTTP 403 by WAF, the correct remediation is typically to identify the triggered rule(s) in the WAF logs and then tune WAF by disabling the specific rule (if justified), creating a targeted exclusion for the specific parameter that causes the false positive, or adjusting the WAF mode (Detection vs Prevention) as appropriate. For controlling access by source IP, you would use WAF custom rules (match condition on RemoteAddr) or Application Gateway features like IP restrictions at the application layer, not a header-value exclusion. See Configure Web Application Firewall on Application Gateway and WAF policy custom rules.

QUESTION NO: 32 - (SIMULATION)

Task 9

You plan to use VNET4 for an Azure API Management implementation.

You need to configure a policy that can be used by an Azure application gateway to protect against known web attack vectors. The policy must only allow requests that originate from IP addresses in Canad

a. You do NOT need to create the application gateway to complete this task.

ANSWER: See the explanation for the answer

Explanation:

To configure a policy in Azure API Management that can be used by an Azure Application Gateway to protect against known web attack vectors and only allow requests from IP addresses in Canada,

follow these steps:

Step-by-Step Solution

Step 1: Create or Access Your API Management Instance Navigate to the Azure Portal.

Search for "API Management services" and select your API Management instance. Step 2: Configure the Policy

In the API Management instance, go to the "APIs" section. Select the API you want to apply the policy to.

Go to the "Design" tab.

Select "All operations" if you want to apply the policy to all operations, or select a specific operation. Step 3: Add the Inbound Policy

In the Inbound processing section, click on "+ Add policy" . Select "IP filter" from the list of policies.

Add the IP address ranges for Canada. You can find the IP ranges for Canada from a reliable source or use a service that provides this information.

Here is an example of the XML configuration for the policy:

Save the policy to apply the changes.

IP Filter Policy: This policy allows you to filter incoming requests based on their IP addresses. By specifying the IP ranges for Canada, you ensure that only requests originating from these IPs are

allowed.

Inbound Processing: Applying the policy in the inbound section ensures that the requests are filtered before they reach your API.

By following these steps, you can configure a policy in Azure API Management that restricts access to your API to only those requests originating from IP addresses in Canada, thereby enhancing security and compliance

QUESTION NO: 33

You have an Azure subscription that contains multiple virtual machines in the West US Azure region. You need to use Traffic Analytics.

Which two resources should you create? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct answer selection is worth one point.

A. an Azure Monitor workbook

B. a Log Analytics workspace

C. a storage account

D. an Azure Sentinel workspace

E. an Azure Monitor data collection rule

ANSWER: B C

Explanation:

To use Traffic Analytics, you must enable NSG flow logs and then have Traffic Analytics process those logs into searchable, aggregated insights. NSG flow logs are written to an Azure Storage account, so creating a storage account is required to store the raw flow log data generated by Network Watcher. Traffic Analytics then reads those flow logs and sends the processed/aggregated results into a Log Analytics workspace, which is where the indexed data is stored and queried to produce the Traffic Analytics dashboards and insights. Without the storage account, there is no destination for the flow logs; without the Log Analytics workspace, Traffic Analytics has nowhere to store and analyze the processed data. These two resources are the core dependencies for enabling Traffic Analytics in a region for your NSGs and virtual machines.

References: Traffic Analytics, NSG flow logs (flow logs overview).

QUESTION NO: 34

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 perform TLS inspection with Azure Firewall for outbound traffic from Azure Virtual Desktop session hosts, you must configure Azure Firewall Premium's TLS inspection feature, which requires the firewall to present a trusted intermediate CA certificate to clients. That certificate is stored and managed in an Azure key vault so the firewall can retrieve and use it for outbound TLS interception. In addition, Azure Firewall needs a way to securely access the key vault without embedding secrets; the recommended approach is to use a managed identity assigned to the firewall and grant it the appropriate Key Vault permissions. With these two resources in place, you can enable TLS inspection on the firewall policy and ensure outbound flows from the host pool are routed through the firewall (typically via UDR), allowing the firewall to decrypt, inspect, and re-encrypt TLS traffic using the certificate chain from Key Vault. This aligns with Microsoft's guidance for Azure Firewall Premium TLS inspection and its integration with Key Vault and managed identities.

References: https://learn.microsoft.com/en-us/azure/firewall/premium-features#tls-inspection, https://learn.microsoft.com/en-us/azure/firewall/premium-tls-inspection

QUESTION NO: 35

You have an Azure Front Door instance that has a single frontend named Frontend1 and an Azure Web Application Firewall (WAF) policy named Policy1. Policy1 redirects requests that have a header containing "string1" to https://www.contoso.com/redirect1. Policy1 is associated to Frontend1.

You need to configure additional redirection settings. Requests to Frontend1 that have a header containing "string2" must be redirected to https://www.contoso.com/redirect2.

Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Create a custom rule.

B. Configure a managed rule.

C. Create a frontend host.

D. Create a policy.

E. Create an association.

F. Add a custom rule to Policy1.

ANSWER: A E F

Explanation:

To redirect requests based on an HTTP request header in Azure Front Door WAF, you use WAF custom rules. Custom rules let you define match conditions (such as a specific request header containing a value) and then take an action. For redirection scenarios, the action is implemented via a custom rule that matches the header condition and returns a redirect response to the specified URL. Because the requirement is to add another redirect behavior to the same frontend, you should add an additional custom rule to the existing WAF policy so that the policy contains both header-based redirect rules (one for "string1" and one for "string2"). After the rule exists in the policy, the policy must be associated with the relevant Front Door frontend (or remain associated if already), ensuring the policy is applied to requests hitting that frontend. This aligns with how Front Door WAF policies are applied: rules live inside a policy, and the policy is bound to a frontend endpoint via an association. For details on WAF custom rules and policy associations in Front Door, see https://learn.microsoft.com/en-us/azure/web-application-firewall/afds/waf-custom-rules and https://learn.microsoft.com/en-us/azure/web-application-firewall/afds/afds-overview.

QUESTION NO: 36

You have an Azure subscription that contains multiple virtual machines in the West US Azure region.

You need to use Traffic Analytics.

Which two resources should you create? Each correct answer presents part of the solution. (Choose two.) NOTE: Each correct answer selection is worth one point.

A. an Azure Monitor workbook

B. a Log Analytics workspace

C. a storage account

D. an Azure Sentinel workspace

E. an Azure Monitor data collection rule

ANSWER: B C

Explanation:

Traffic Analytics is a Network Watcher capability that analyzes NSG flow logs to provide visibility into traffic patterns, top talkers, and security insights. To use it, you need a place to store the raw NSG flow logs and a place to ingest and query the processed telemetry. NSG flow logs are written to an Azure Storage account, which is required as the log destination. Traffic Analytics then processes those flow logs and sends the enriched results to a Log Analytics workspace, which is required for the analytics, querying, and built-in visualizations. Once the Log Analytics workspace is connected, you can view Traffic Analytics insights through the Azure portal experience that reads from that workspace. These two resources are the foundational dependencies for enabling NSG flow logs with Traffic Analytics across your virtual machines' subnets/NSGs in the region.

References: Traffic Analytics (Network Watcher), NSG flow logs overview

QUESTION NO: 37

You plan to deploy an Azure virtual network. You need to design the subnets.

Which three types of resources require a dedicated subnet? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. VPN gateway

B. Azure Bastion

C. Azure Active Directory Domain Services (Azure AD DS)

D. Azure Application Gateway v2

E. Azure Private Link

ANSWER: A B D

Explanation:

Several Azure networking services must be deployed into their own dedicated subnet because the platform enforces specific subnet names, applies service-managed network policies, and requires exclusive control of IP addressing and routing within that subnet. A VPN gateway requires a dedicated subnet named **GatewaySubnet**; Azure uses this subnet to place gateway instances and manage routing for site-to-site, point-to-site, and VNet-to-VNet connectivity. Azure Bastion also requires a dedicated subnet named **AzureBastionSubnet** so the managed Bastion service can provide secure RDP/SSH access without exposing VMs to public IPs. Azure Application Gateway v2 must be deployed into its own subnet as well, because the gateway instances and their scaling/updates are managed within that subnet and it can't share the subnet with other resource types. These requirements are core VNet design considerations to avoid deployment failures and to ensure the services can scale and be maintained by Azure. See Microsoft guidance on dedicated subnets for these services at https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-for-azure-services and Application Gateway subnet requirements at https://learn.microsoft.com/en-us/azure/application-gateway/configuration-infrastructure.

QUESTION NO: 38 - (HOTSPOT)

HOTSPOT

You have an Azure subscription that contains the route tables and routes shown in the following table.

Route table name	Route name	Prefix	Destination
RT1	Default Route	0.0.0.0/0	VirtualNetworkGateway
RT2	Default Route	0.0.0.0/0	Internet

The subscription contains the subnets shown in the following table.

Name	Prefix	Route table	Virtual network
Subnet1	10.10.1.0/24	RT1	Vnet1
Subnet2	10.10.2.0/24	RT2	Vnet1
GatewaySubnet	10.10.3.0/24	**None**	Vnet1

The subscription contains the virtual machines shown in the following table.

Name	IP address
VM1	10.10.1.5
VM2	10.10.2.5

There is a Site-to-Site VPN connection to each local network gateway.

For each of the following statements, select Yes of the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Traffic from VM2 to the internet is routed through the New-York Site-to-Site VPN connection	◯	◯
Traffic from VM1 to VM2 is routed through the New-York Site-to-Site VPN connection	◯	◯
Traffic from VM1 to the internet is routed through the New-York Site-to-Site VPN connection	◯	◯

ANSWER:

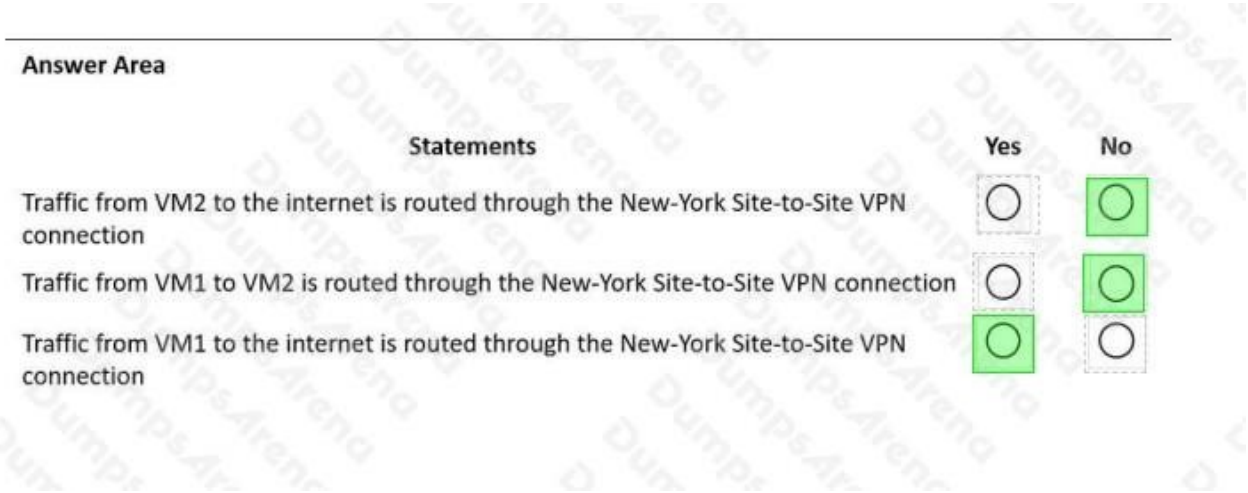

Explanation:

This question tests your understanding of Azure user-defined routes (UDRs) and how they affect traffic routing in virtual networks. Let's break down each statement by analyzing the routing configuration.

Understanding the Setup: We have two route tables (RT1 and RT2), both containing default routes (0.0.0.0/0) but with different next hops. RT1 points to VirtualNetworkGateway (forcing traffic through the Site-to-Site VPN), while RT2 points directly to Internet. VM1 (10.10.1.5) is in Subnet1 which uses RT1, and VM2 (10.10.2.5) is in Subnet2 which uses RT2. Both VMs are in the same virtual network (Vnet1).

Statement 1 - VM2 to Internet (NO): VM2 resides in Subnet2, which is associated with route table RT2. When VM2 sends traffic to the internet, Azure looks up the effective routes. RT2 contains a default route (0.0.0.0/0) with next hop type "Internet", meaning traffic goes directly to the internet without passing through any VPN gateway. Therefore, VM2's internet traffic does NOT route through the New-York Site-to-Site VPN connection. The answer is **NO**.

Statement 2 - VM1 to VM2 (NO): This is intra-VNet communication. Even though VM1's subnet (Subnet1) has RT1 associated with it, which contains a route to VirtualNetworkGateway, Azure's routing logic prioritizes more specific routes. When VM1 (10.10.1.5) sends traffic to VM2 (10.10.2.5), the destination address falls within the VNet's address space. Azure automatically creates system routes for VNet address spaces, and these are more specific than the 0.0.0.0/0 catch-all route. According to the longest prefix match rule, the VNet system route wins, so traffic between VM1 and VM2 stays within the virtual network and does NOT traverse the VPN gateway. The answer is **NO**.

Statement 3 - VM1 to Internet (YES): VM1 is in Subnet1, which uses route table RT1. RT1 contains a default route (0.0.0.0/0) with next hop type "VirtualNetworkGateway". This configuration is known as forced tunneling. When VM1 attempts to reach any internet destination, the 0.0.0.0/0 route matches all internet-bound traffic and directs it to the virtual network gateway (the Site-to-Site VPN connection to New York). This means all of VM1's internet traffic is routed through the VPN connection. The answer is **YES**.

Key Concept: The critical principle here is understanding how Azure selects routes. Azure uses longest prefix matching, meaning more specific routes take precedence. For internet traffic, if a UDR with 0.0.0.0/0 pointing to VirtualNetworkGateway exists in the associated route table, forced tunneling occurs. For intra-VNet traffic, the default VNet system routes (which are more specific than 0.0.0.0/0) ensure traffic stays local unless explicitly overridden by an even more specific UDR.

For more details on Azure routing behavior, see the official documentation on User-defined routes and how Azure selects a route.

QUESTION NO: 39

You plan to deploy Azure virtual network.

You need to design the subnets.

Which three types of resources require a dedicated subnet? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Azure Bastion

B. Azure Active Directory Domain Services

C. Azure Private Link

D. Azure Application Gateway v2

E. VPN gateway

ANSWER: A D E

Explanation:

Several Azure networking services require a dedicated subnet because Azure enforces reserved subnet names and/or requires exclusive control of IPs and network policies for the service to function correctly. Azure Bastion must be deployed into a subnet named **AzureBastionSubnet** that is dedicated to Bastion; this is required so the managed Bastion service can scale and operate securely within your virtual network. Azure Application Gateway v2 also requires its own dedicated subnet because the gateway instances are deployed into that subnet and need exclusive IP capacity and configuration control for scaling and routing. Similarly, a VPN gateway must be deployed into a dedicated subnet named **GatewaySubnet**; Azure uses this subnet to place gateway instances and manage routing and tunnel endpoints, and it must not contain other resource types. These requirements are part of Azure's virtual network design guidance for platform-managed networking services and are commonly tested in AZ-700 because they affect address planning and subnet layout.

References: https://learn.microsoft.com/en-us/azure/bastion/bastion-overview, https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-about-vpngateways

QUESTION NO: 40

You need to use Traffic Analytics to monitor the usage of applications deployed to Azure virtual machines.

Which Azure Network Watcher feature should you implement first?

A. Connection monitor

B. Packet capture

C. NSG flow logs

D. IP flow verify

ANSWER: C

Explanation:

To use Traffic Analytics, you must first enable collection of network traffic data from your workloads. Traffic Analytics is built on top of NSG flow logs: it analyzes the flow records produced for network interfaces that are associated with a Network Security Group and then visualizes insights such as traffic distribution, top talkers, and communication patterns. Without NSG flow logs enabled (and configured to send logs to a storage account and/or Log Analytics workspace), Traffic Analytics has no underlying flow data to process, so it can't provide application usage and traffic insights for your Azure virtual machines. Therefore, the first Azure Network Watcher feature to implement is NSG flow logs, and then you enable Traffic Analytics on those flow logs to start generating analytics in Log Analytics.

References: Microsoft Docs: Traffic Analytics, Microsoft Docs: NSG flow logs overview

QUESTION NO: 41

You have an Azure subscription that contains multiple virtual machines in the West US Azure region.

You need to use Traffic Analytics.

Which two resources should you create? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct answer selection is worth one point.

A. an Azure Monitor workbook

B. a Log Analytics workspace C a storage account

C. an Azure Sentinel workspace

D. an Azure Monitor data collection rule

E. a storage account

ANSWER: B E

Explanation:

To use Traffic Analytics, you must enable NSG flow logs and configure Traffic Analytics processing. NSG flow logs are written to an Azure Storage account, which is a required dependency because the raw flow log data is stored there before it's processed. Traffic Analytics then processes those flow logs and sends the enriched, aggregated results into a Log Analytics workspace, where the data is indexed and becomes available for querying and visualization. In practice, the workflow is: create (or select) a Storage account for NSG flow logs, create (or select) a Log Analytics workspace, then enable NSG flow logs and turn on Traffic Analytics to point at that workspace. These two resources are the core platform components Traffic Analytics relies on; other items like workbooks can be used for visualization but aren't required resources to "use Traffic Analytics." For implementation details and prerequisites, see Microsoft's Traffic Analytics documentation and the steps for enabling flow logs and Traffic Analytics settings.

Microsoft Docs: Traffic Analytics
Microsoft Docs: NSG flow logs overview

QUESTION NO: 42

You need to configure APPGW1 to support end-to-end encryption. The solution must meet the security requirements. What should you do?

A. From the SSL settings, upload a TLS client certificate that is issued by the internal root CA and includes the full certificate chain.

B. From the Backend settings, upload a wildcard TLS certificate that has a private key issued by the internal root CA

C. From the Backend settings, upload the internal root CA certificate.

D. From the SSL settings, upload a TLS client certificate that is issued by the internal root CA.

ANSWER: C

Explanation:

To support end-to-end encryption with Azure Application Gateway, you must configure the gateway to use HTTPS from the gateway to the backend targets and ensure the gateway can validate the backend server certificate. This is done by uploading the trusted root certificate (or intermediate chain as needed) that issued the backend's TLS certificate into the backend HTTP settings (called "Backend settings" in the portal). Application Gateway uses this trusted root certificate to authenticate the backend during the TLS handshake and prevent man-in-the-middle attacks. This aligns with typical security requirements for end-to-end TLS: encrypt traffic all the way to the backend and validate the backend identity using a trusted CA. Uploading a client certificate is only required when you enable mutual TLS (mTLS) to the backend, which is a different requirement than standard end-to-end encryption. Likewise, uploading a server certificate with a private key is used for the listener (front-end) termination, not for establishing trust to the backend. Therefore, the correct action is to upload the internal root CA certificate in the Backend settings so the gateway trusts the backend certificate chain.

References: https://learn.microsoft.com/en-us/azure/application-gateway/ssl-overview, https://learn.microsoft.com/en-us/azure/application-gateway/configuration-http-settings