PassQueen
a.k.a dumpsqueen.com

Designing and Implementing Microsoft Azure Networking Solutions

Microsoft AZ-700

Version Demo

Total Demo Questions: 52

Total Premium Questions: 527

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Design, implement, and manage hybrid networking	109
Topic 2, Design and implement core Azure networking infrastructure	167
Topic 3, Design and implement routing	68
Topic 4, Secure and monitor networks	146
Topic 5, New Topic: Contoso Case Study	17
Topic 6, New Topic: Proseware. Inc Case Study	15
Topic 7, Case Study 1	2
Topic 8, Case Study 2	1
Topic 9, Case Study 3	2
Total	527

QUESTION NO: 1

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

You need to implement a Site-to-Site VPN connection between the datacenter and the virtual network.

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

A. a virtual network gateway

B. Azure Firewall

C. a local network gateway

D. Azure Web Application Firewall (WAF)

E. an on-premises data gateway

F. an Azure application gateway

G. a user-defined route

ANSWER: A C

Explanation:

A Site-to-Site VPN requires an Azure VPN Gateway deployment that has definitions for both ends of the IPsec/IKE tunnel. A virtual network gateway is the Azure-side gateway resource. It is deployed to the dedicated `GatewaySubnet` in the virtual network, has a public IP address, and terminates VPN traffic for connections to external networks. The gateway must use a VPN gateway SKU and configuration appropriate for the required connectivity and throughput.

A local network gateway represents the on-premises side of the connection in Azure. Its configuration identifies the public IP address (or fully qualified domain name) of the on-premises VPN device and includes the on-premises address prefixes that Azure must route through that device. After both gateway resources exist, a connection resource associates the virtual network gateway with the local network gateway and configures the shared key and, when required, IPsec/IKE policy settings. This resource model enables bidirectional routing between the Azure virtual network and the declared on-premises networks. Microsoft's deployment guidance describes creating the virtual network gateway, local network gateway, and connection as the standard Site-to-Site configuration flow. See Create a Site-to-Site VPN connection and About Azure VPN Gateway.

QUESTION NO: 2

You plan to publish a website that will use an FQDN of www.contoso.com. The website will be hosted by using the Azure App Service apps shown in the following table.

Name	FQDN	Location	Public IP address
AS1	As1.contoso.com	East US	131.107.100.1
AS2	As2.contoso.com	West US	131.107.200.1

You plan to use Azure Traffic Manager to manage the routing of traffic for www.contoso.com between AS1 and AS2.

You need to ensure that Traffic Manager routes traffic for www.contoso.com. Which DNS record should you create?

A. two A records that map wmv.contoso.com to 131 107 100 1 and 131 107 200 1

B. a CNAME record that maps www.contoso.com to TMprofile1.azurefd.net

C. a CNAME record that maps www.contoso.com to TMprofile1.trafficmanager.net

D. a TXT record that contains a string ofas1.contoso.com and as2.contoso.com in the details

ANSWER: C

Explanation:

A CNAME record that maps www.contoso.com to TMprofile1.trafficmanager.net is correct because Azure Traffic Manager is a DNS-based traffic-routing service. A Traffic Manager profile receives a DNS name in the `trafficmanager.net` domain, and a custom subdomain such as `www.contoso.com` must alias that profile name through a CNAME record. When a client resolves the `www` hostname, DNS resolution reaches the Traffic Manager profile. Traffic Manager then evaluates its configured routing method and endpoint health, and returns the DNS response for the appropriate Azure App Service endpoint. This enables the profile to direct requests between AS1 and AS2 and to respond to endpoint availability changes without requiring clients to know either app's direct hostname.

Using a CNAME is appropriate because `www.contoso.com` is a subdomain rather than the zone apex. Each App Service app should also be configured to recognize the custom hostname so that it can serve requests addressed to `www.contoso.com` after Traffic Manager routes them. Microsoft documents this custom-domain pattern for Traffic Manager and App Service at Configure a custom domain with Traffic Manager and describes Traffic Manager's DNS-based routing behavior at Azure Traffic Manager overview.

QUESTION NO: 3

You have an Azure application gateway named AppGW1 that balances requests to a web app named App1.

You need to modify the server variables in the response header of App1. What should you configure on AppGW1?

A. HTTP settings

B. rewrites

C. rules

D. listeners

ANSWER: B

Explanation:

rewrites is correct because Azure Application Gateway uses rewrite sets to modify HTTP request and response headers as traffic passes through the gateway. A rewrite set contains one or more rewrite rules, and each rule can perform actions such as setting, adding, or deleting a response-header value before Application Gateway returns the backend response to the client. Rewrite rules can also use server variables, which expose information about the request, client, listener, and gateway processing context. For example, a response-header action can insert a server-variable value into a header, or alter an existing header based on gateway-provided context. After creating the rewrite set, associate it with the applicable routing rule or path-based routing configuration so it is applied to requests sent to App1. This enables response-header modification at the Application Gateway layer and does not require changing the web application's code or its web server configuration. Microsoft documents header and URL rewriting, including the use of server variables in rewrite actions, in Rewrite HTTP headers and URL with Application Gateway. For the relationships among gateway routing components and rules, see Application Gateway configuration overview.

QUESTION NO: 4

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). You configure a service tag for Microsoft.Storage and link the tag to Subnet1.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. The `Microsoft.Storage` service tag represents Azure Storage public endpoint address ranges, rather than an individual storage account such as `storage1`. Applying an NSG to Subnet1 and using this service tag in an outbound rule can allow VM1 to communicate with Azure Storage endpoints covered by the tag, but it cannot distinguish `storage1` from `storage2` or `storage3`. Consequently, it does not enforce the required per-storage-account restriction.

Service tags are designed to simplify NSG rule maintenance by representing groups of Azure service IP prefixes. They are not resource-specific access controls. To grant the VM access to only `storage1`, the restriction must be enforced at the storage-account level. For example, configure `storage1` to allow the appropriate virtual network subnet through its networking settings, using a virtual network service endpoint, or expose `storage1` through a private endpoint. Keep public network access restricted on the other storage accounts so that the subnet used by VM1 is not authorized to reach them. This creates an access boundary specific to the intended storage account. See Azure service tags overview and Configure Azure Storage firewalls and virtual networks.

QUESTION NO: 5

You have a virtual network named Vnet1 that contains a virtual machine named VM1. VM1 cannot connect to a private IP address of 10.20.1.4 on TCP port 443.

You need to identify whether Azure routing or a network security group blocks the connection. The solution must use Azure Network Watcher.

Which two features should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Next hop

B. IP flow verify

C. Topology

D. Packet capture

E. Connection Monitor

ANSWER: A B

Explanation:

Next hop identifies the route that Azure selects for traffic from a virtual machine network interface to a specified destination IP address. It reports the next-hop type and, when applicable, the next-hop IP address. This helps determine whether a user-defined route, virtual network gateway, virtual appliance, or another route is affecting the connection.

IP flow verify evaluates whether a specified TCP or UDP flow is allowed or denied by the network security groups applied to the virtual machine network interface and subnet. When the flow is denied, the result identifies the NSG rule responsible. Together, these features isolate routing and NSG causes without requiring packet capture. See Next hop overview and IP flow verify overview.

QUESTION NO: 6

You need to manage connectivity from NYCNet to the Azure services that use private endpoints. The solution must meet the security requirements. What should you do first?

A. Add a route table to SUBNET-PL

B. Enable private endpoint network policies for SUBNET-PE.

C. From Azure Virtual Network Manager, create a security admin configuration.

D. From Azure Virtual Network Manager, create a network group that has Member type set to Subnet

ANSWER: B

Explanation:

Enable private endpoint network policies for SUBNET-PE first. A private endpoint is assigned a network interface and private IP address in its delegated subnet, but Azure disables private endpoint network policies on that subnet by default. Enabling the policy removes that default limitation and allows network security groups and user-defined routes to be evaluated for traffic to and from private endpoints. This provides the required foundation for applying centralized network-security controls to connectivity between NYCNet and services exposed through private endpoints.

In particular, enabling the policy is necessary before relying on subnet-level traffic control and routing behavior to enforce the organization's security design. The policy can be enabled for all private endpoints in the subnet through the subnet's `privateEndpointNetworkPolicies` setting. It should be enabled before implementing the higher-level management configuration so that the private endpoint subnet can participate in the intended security enforcement model. Microsoft documents the private endpoint network-policy setting and its support for NSGs and UDRs in Manage network policies for private endpoints. For the role of centrally managed security rules in Azure Virtual Network Manager, see Security admin rules in Azure Virtual Network Manager.

QUESTION NO: 7 - (DRAG DROP)

Your on-premises network contains two subnets named Subnet1 and Subnet2. Subnet2 contains a Hyper-V host that contains two virtual machines named VM1 and VM2. VM1 and VM2 are connected to Subnet2.

You have an Azure virtual network named VNet1 that contains GatewaySubnet and a subnet named VSubnet1. VNet1 is connected to the on-premises network by using a Site-to-Site (S2S) VPN connection.

You plan to migrate VM1 to VNet1 and maintain the existing IP address of VM1. VM2 will remain on Subnet2.

You need to prepare the environment to ensure that VM1 can communicate with VM2 once the migration is complete.

Which five 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:

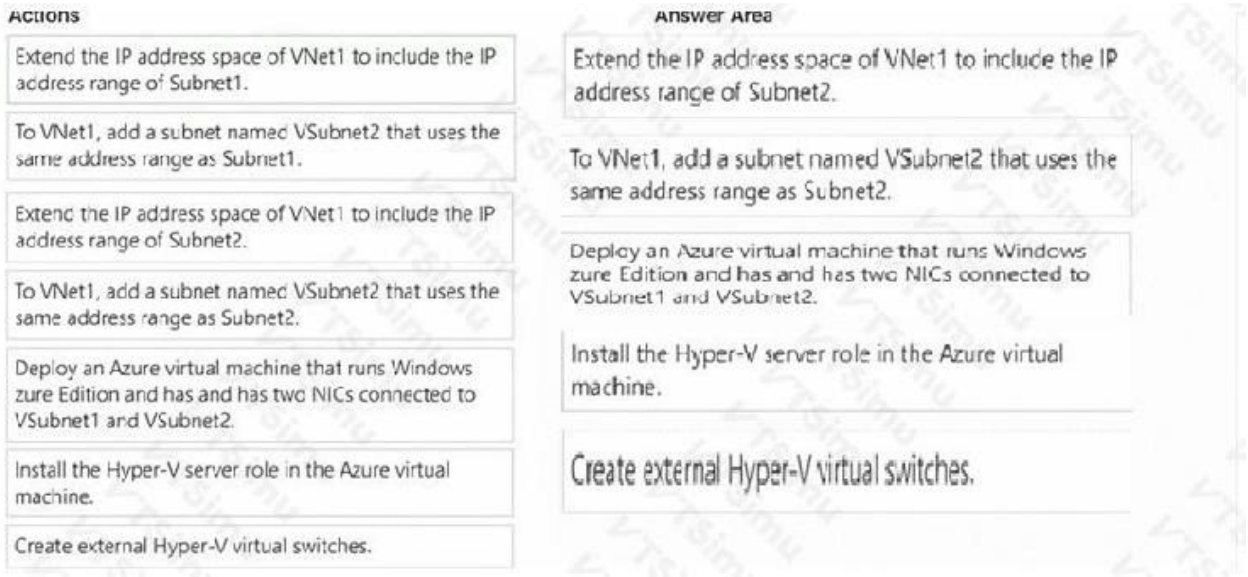

Explanation:

The required sequence prepares Azure to host the migrated workload while preserving VM1's existing address from Subnet2. First, VNet1 must be expanded to include the address range used by Subnet2. This makes the required address space available in the Azure virtual network configuration. Next, VSubnet2 is created from that address range. The additional subnet provides the Azure-side network attachment needed for the networking design that preserves the workload's existing IP addressing during the migration.

A Windows Server Azure Edition virtual machine is then deployed with one network interface connected to VSubnet1 and another connected to VSubnet2. Windows Server Azure Edition supports nested virtualization scenarios in Azure, allowing this Azure VM to act as the Hyper-V host for the migrated virtual machine. The two NIC design supplies the necessary network connections for the Hyper-V host and its virtual networking configuration. Microsoft documents the Azure Edition capabilities and nested virtualization prerequisites in the Windows Server Azure Edition overview and the nested virtualization guidance for Azure VMs.

After the host VM exists, the Hyper-V server role is installed. Installing the role enables the Azure VM to run VM1 as a nested virtual machine, which is the key hosting component for this migration approach. Finally, external Hyper-V virtual switches are created and associated with the host's network adapters. These switches provide the nested VM with the required connectivity through the Azure host's NICs, allowing VM1 to use its retained network identity and communicate with VM2 across the established site-to-site VPN connectivity. Azure VPN Gateway provides the private cross-premises connection used by this design, as described in the Azure VPN Gateway documentation.

QUESTION NO: 8

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.

You need to ensure that the URL is accessible through the application gateway. Solution: You add a rewrite rule for the host header.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. An Application Gateway rewrite rule cannot resolve a WAF-generated HTTP 403 caused by host-header validation. Application Gateway does not support using rewrite rules to modify the `Host` header. More importantly, WAF evaluates the incoming request as it reaches the gateway, so a request that WAF blocks based on its original host-header characteristics is denied before a rewrite rule could make it acceptable for this purpose.

The appropriate approach is to send a request containing a valid host name that is configured on an Application Gateway listener. For a custom application domain, configure a multi-site listener with that domain name and ensure the corresponding DNS record resolves to the Application Gateway frontend IP address. If the requirement concerns the host name sent onward to a backend service, configure the backend HTTP settings to preserve the incoming host name or override it with the required backend domain name. These settings handle backend host-header behavior and do not attempt to bypass WAF inspection of an invalid incoming request.

Microsoft documents the rewrite-rule restrictions, including the limitation for the `Host` header, at Application Gateway rewrite limitations. Listener host-name configuration is described in Hosting multiple sites with Application Gateway.

QUESTION NO: 9

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:

A virtual network gateway and a local network gateway are required Azure resources for a Site-to-Site VPN connection. The virtual network gateway is the Azure-side VPN endpoint. It is deployed in a dedicated subnet named `GatewaySubnet` within the target virtual network and provides the VPN Gateway service that terminates IPsec/IKE tunnels. Its configuration includes a gateway type of VPN, an appropriate SKU, and a public IP address resource used by the on-premises VPN device to establish the tunnel.

A local network gateway represents the on-premises side of the connection within Azure. It contains the public IP address or fully qualified domain name of the on-premises VPN device, along with the address prefixes of the on-premises networks that Azure should route through that device. Azure then uses these two gateway resources when a VPN connection resource is created, with matching shared-key and IPsec/IKE settings configured on the on-premises VPN device.

These resources provide the endpoint and routing information required to model and establish the cross-premises VPN relationship. For implementation guidance, see Create a site-to-site VPN connection in the Azure portal and About VPN Gateway.

QUESTION NO: 10

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.private.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:

Create a private endpoint for webapp1 is required because an Azure App Service private endpoint assigns a private IP address from the virtual network to the App Service. Virtual machines in Vnet1 can then send traffic to the app through Azure Private Link, without requiring outbound Internet connectivity. The private endpoint is the network path that makes private, inbound access to the multitenant App Service possible.

Create a CNAME record that maps www.private.contoso.com to webapp1.privatelink.azurewebsites.net supplies the required name-resolution path for the requested custom hostname. The CNAME directs requests for the custom name to the App Service Private Link hostname. DNS for `privatelink.azurewebsites.net` must resolve the app hostname to the private endpoint IP address for clients in Vnet1; this is ordinarily implemented by linking the appropriate private DNS zone to Vnet1 and ensuring it contains the endpoint record. As a result, a VM resolves the custom URL to the private endpoint and reaches webapp1 entirely through private connectivity. App Service must also have the custom hostname configured as a domain binding and an appropriate TLS certificate to serve HTTPS for that hostname. See Use private endpoints for Azure App Service and Azure Private Endpoint DNS configuration.

QUESTION NO: 11

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

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

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

D. a real user measurements key in Traffic Manager

ANSWER: A B

Explanation:

A Traffic Manager profile that uses the priority routing method provides the required active-passive failover behavior. Priority routing directs all traffic to the healthy endpoint assigned the highest priority, so App1eu can be configured as the primary endpoint. Traffic Manager continuously evaluates endpoint health by using its configured health probes. If App1eu does not respond successfully to those probes, Traffic Manager removes it from DNS responses and returns App1us, which is configured with a lower priority, instead. When App1eu becomes healthy again, it resumes receiving traffic because it remains the preferred endpoint.

A CNAME record in a DNS domain named fabrikam.com is also required to make the custom host name resolve through Traffic Manager. The `www` host record should be configured as a CNAME whose target is the DNS name assigned to the Traffic Manager profile, such as `<profile>.trafficmanager.net`. Clients then resolve `www.fabrikam.com` through Traffic Manager and receive the currently appropriate application endpoint based on priority and health. Microsoft documents priority routing and endpoint failover at Traffic Manager routing methods, and documents custom-domain CNAME configuration at Point an internet domain to an Azure Traffic Manager profile.

QUESTION NO: 12

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 depends on network flow-log data being retained and then analyzed. A storage account is required because Network Watcher flow logs are written to Azure Storage. This account provides the destination for the raw log files generated from the monitored network security groups and their associated network interfaces. The storage destination must be available in the appropriate Azure region when configuring the flow-log collection.

A Log Analytics workspace is also required because Traffic Analytics processes the flow logs and sends its enriched, aggregated analytics results to that workspace. The workspace makes the data available for Azure Monitor Logs queries, visualizations, and Traffic Analytics insights, such as traffic patterns, allowed and denied flows, and communications between workloads. Therefore, creating both a storage account and a Log Analytics workspace establishes the data-storage and analytics destinations needed for Traffic Analytics.

Microsoft documents the Traffic Analytics configuration and its Log Analytics workspace requirement at Traffic Analytics. The role of Azure Storage as the flow-log destination is described in the virtual network flow logs overview.

QUESTION NO: 13

You plan to publish a website that will use an FQDN of www.contoso.com. The website will be hosted by using the Azure App Service apps shown in the following table.

Name	FQDN	Location	Public IP address
AS1	As1.contoso.com	East US	131.107.100.1
AS2	As2.contoso.com	West US	131.107.200.1

You plan to use Azure Traffic Manager to manage the routing of traffic for www.contoso.com between AS1 and AS2.

You need to ensure that Traffic Manager routes traffic for www.contoso.com.

Which DNS record should you create?

A. two A records that map www.contoso.com to 131.107.100.1 and 131.107.200.1

B. a CNAME record that maps www.contoso.com to TMprofile1.azurefd.net

C. a CNAME record that maps www.contoso.com to TMprofile1.trafficmanager.net

D. a TXT record that contains a string of as1.contoso.com and as2.contoso.com in the details

ANSWER: C

Explanation:

A CNAME record that maps www.contoso.com to TMprofile1.trafficmanager.net is required because Azure Traffic Manager operates at the DNS layer. Traffic Manager does not act as an inline reverse proxy for web requests; instead, a client first queries DNS for www.contoso.com. The CNAME alias causes that query to resolve through the Traffic Manager profile's trafficmanager.net name. Traffic Manager then evaluates the profile's configured routing method and endpoint health and returns the DNS response for the appropriate App Service endpoint.

This configuration lets Traffic Manager control routing between AS1 and AS2 while retaining the public custom hostname, www.contoso.com. It also avoids coupling the public DNS record to individual endpoint IP addresses, which can change. Each App Service app must additionally be configured to recognize the www.contoso.com custom hostname so that it can serve requests containing that host header. Microsoft's guidance for using a custom domain with App Service and Traffic Manager specifically directs you to create a CNAME record from the custom subdomain to the Traffic Manager profile DNS name. See Azure Traffic Manager overview and Configure a custom domain name for App Service with Traffic Manager.

QUESTION NO: 14 - (HOTSPOT)

You need to configure connectivity between NYCNet and SFONet. The solution must meet the connectivity requirements. What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

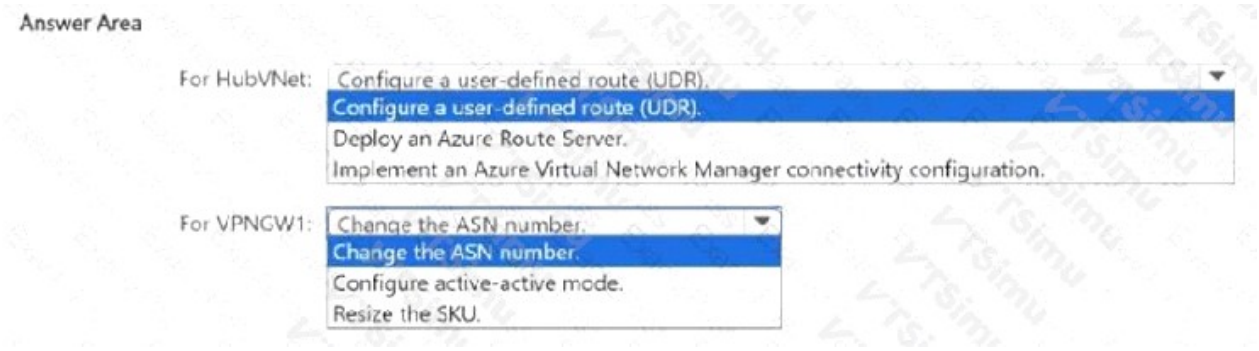

ANSWER:

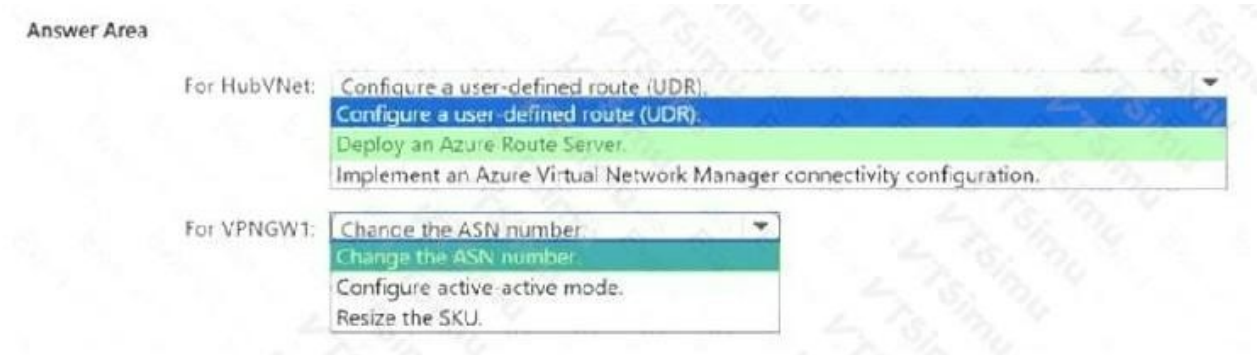

Explanation:

Deploy Azure Route Server in HubVNet and change the ASN number on VPNGW1. Azure Route Server is the managed routing service that enables dynamic route exchange by using BGP between supported network appliances and Azure VPN Gateway. In a hub design, this gives the hub a centralized BGP route-exchange point, allowing learned prefixes to be

distributed automatically as connectivity information changes. This is the appropriate mechanism when the required connectivity depends on dynamic routing rather than individually maintained static routes.

The VPN gateway must use a distinct BGP autonomous system number from the Route Server. Azure Route Server uses an ASN for its BGP sessions, and a BGP peer cannot establish the required peering relationship when both sides use the same ASN. Changing VPNGW1's ASN ensures that the gateway can form BGP peering with the Route Server and participate in route advertisement and learning. Once this peering is established, the relevant network prefixes can be exchanged dynamically, supporting the required communication path between NYCNet and SFONet through the hub.

Azure Route Server is deployed in a dedicated subnet named `RouteServerSubnet`, and the VPN gateway remains deployed in its required gateway subnet. The Route Server does not replace VPNGW1; instead, it complements the gateway by providing managed BGP route exchange. Microsoft documents the Route Server architecture and BGP integration at Azure Route Server overview. The configuration requirements for using Route Server with Azure VPN Gateway, including the need for distinct ASN values, are described at Azure Route Server support for ExpressRoute and VPN Gateway. This combination provides the dynamic, scalable routing foundation needed for the connectivity design.

QUESTION NO: 15

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

Name	Type	Description
FW1	Azure Firewall Premium	Has a network intrusion detection and prevention system (IDPS) enabled
HP1	Azure Virtual Desktop host pool	All outbound traffic from HP1 to the subscription's resources route through FW1
Server1	Virtual machine	Hosts an application named App1
KV1	Azure Key Vault	*None*

Users on HP1 connect to App1 by using a URL of https://app1 .comoso.com.

You need to ensure that the IDPS on FW1 can identify security threats in the connections from HP1 to Server1.

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

A. Enable TLS inspection for FW1.

B. Import an intermediate CA certificate to KV1.

C. Enable threat intelligence for FW1.

D. Add an application group to HP1.

E. Add a secured virtual network to FW1.

ANSWER: A B

Explanation:

Enable TLS inspection for FW1 and import an intermediate CA certificate to KV1. The connection uses HTTPS, which encrypts the application-layer traffic exchanged between HP1 and the application. Azure Firewall Premium IDPS can inspect encrypted traffic payloads only when TLS inspection is enabled. TLS inspection causes the firewall to terminate the client TLS session, decrypt the traffic for IDPS and other security-policy evaluation, and establish a separate TLS session to the destination. This lets IDPS evaluate the HTTP content and identify signatures or malicious activity that would otherwise remain encrypted.

Azure Firewall TLS inspection requires a customer-managed intermediate certificate authority certificate. The certificate, including its private key, is stored in Azure Key Vault and referenced by the firewall policy TLS inspection configuration. The firewall uses that intermediate CA certificate to generate certificates presented to clients during inspection. For HP1 to trust the inspection process, the corresponding root CA must also be trusted on the client. With the intermediate CA certificate

available in KV1 and TLS inspection enabled, FW1 can decrypt and inspect the HTTPS flow to Server1. See Azure Firewall Premium TLS inspection and Configure TLS inspection for Azure Firewall Premium.

QUESTION NO: 16

You have an Azure virtual network that contains the subnets shown in the following table.

You deploy an Azure firewall to AzureFirewallSubnet. You route all traffic from Subnet2 through the firewall.

You need to ensure that all the hosts on Subnet2 can access an external site located at https://*.contoso.com.

What should you do?

A. Create a network security group (NSG) and associate the NSG to Subnet2.

B. In a firewall policy, create an application rule.

C. In a firewall policy, create a DNAT rule.

D. In a firewall policy, create a network rule.

ANSWER: B

Explanation:

In a firewall policy, create an application rule. Azure Firewall application rules are designed for outbound HTTP and HTTPS traffic and can use fully qualified domain names (FQDNs), including wildcard FQDNs such as `*.contoso.com`. Configure an allow rule collection that applies to the relevant source, such as Subnet2 or its IP address range, select the HTTPS protocol, and specify `*.contoso.com` as the target FQDN. Azure Firewall then evaluates the requested web destination and permits matching HTTPS connections through the firewall. This is the appropriate control because the access requirement is expressed as a web protocol and domain-name pattern rather than only as an IP address, port, or inbound translation requirement. Application rules support explicit HTTP/S filtering and Azure Firewall performs the required processing for the outbound web session. Ensure the rule collection is assigned to the firewall policy associated with the deployed Azure Firewall and has a priority that allows the intended traffic. Microsoft documents application rules and their FQDN target support in Azure Firewall application rules. For rule-processing behavior and policy configuration, see Azure Firewall rule processing.

QUESTION NO: 17

Your company has offices in New York and Amsterdam. The company has an Azure subscription. Both offices connect to Azure by using a Site-to-Site VPN connection.

The office in Amsterdam uses resources in the North Europe Azure region. The office in New York uses resources in the East US Azure region.

You need to implement ExpressRoute circuits to connect each office to the nearest Azure region. Once the ExpressRoute circuits are connected, the on-premises computers in the Amsterdam office must be able to connect to the on-premises servers in the New York office by using the ExpressRoute circuits.

Which ExpressRoute option should you use?

A. ExpressRoute Local

B. ExpressRoute FastPath

C. ExpressRoute Direct

D. ExpressRoute Global Reach

ANSWER: D

Explanation:

ExpressRoute Global Reach is the appropriate feature because it provides private, on-premises-to-on-premises connectivity by linking two ExpressRoute circuits across the Microsoft global network. Here, the Amsterdam office can connect through an ExpressRoute circuit in North Europe, while the New York office uses another circuit in East US. Global Reach connects the Azure private peering associated with those circuits, allowing routes between the two respective on-premises networks to be exchanged and carried over Microsoft's backbone.

This enables computers in Amsterdam to access servers in New York without using the public internet or requiring traffic to traverse virtual networks as a transit mechanism. Each site retains its local ExpressRoute connection while Global Reach provides the inter-site WAN path required by the scenario. The connectivity is established by configuring a Global Reach connection between the circuits' private peerings; the circuits must meet the applicable Global Reach requirements, including being in supported peering locations and using the required private peering configuration. See Microsoft Learn: ExpressRoute Global Reach and Microsoft Learn: What is ExpressRoute?.

QUESTION NO: 18

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 ap pear 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.

Sol ution: You implement Azure Firewall.

Does this meet the requirement?

A. Yes

B. No

ANSWER: A

Explanation:

Yes is correct. An Azure Virtual WAN hub becomes a secured virtual hub when Azure Firewall is deployed in that virtual hub. Azure Firewall is the managed network security service integrated with Virtual WAN hubs and provides centralized inspection and policy enforcement for traffic traversing the hub. After the firewall deployment is completed, the hub's security state is shown as secured in the Azure portal. This status indicates that the Virtual WAN hub has an associated security provider, specifically Azure Firewall in this solution; it does not require every possible routing or firewall-policy configuration to be completed merely to change the hub from the unsecured state. Azure Firewall Manager can be used to deploy and manage Azure Firewall in one or more Virtual WAN hubs, including associating a firewall policy and configuring routing intent where needed. Therefore, implementing Azure Firewall in Hub1 directly meets the stated requirement to have Hub1 marked as Secured. See the Microsoft documentation for securing Virtual WAN hubs with Azure Firewall Manager and configuring Azure Firewall in a Virtual WAN hub.

QUESTION NO: 19

You have Azure virtual networks in the East US and West US Azure regions. Each virtual network contains public IP addresses that are associated to virtual machines.

You need to protect the public IP addresses from volumetric and protocol DDoS attacks. The solution must provide DDoS attack telemetry and DDoS Protection cost protection.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure DDoS Protection plan.

B. Associate the virtual networks to the DDoS Protection plan.

C. Enable Just-in-time VM access on the virtual machines.

D. Create a Web Application Firewall (WAF) policy.

E. Associate an Azure Firewall policy to each virtual network.

ANSWER: A B

Explanation:

An **Azure DDoS Protection plan** provides enhanced DDoS mitigation capabilities, telemetry, and cost protection for protected resources. A single plan can protect public IP resources in multiple virtual networks, including virtual networks in different regions.

After creating the plan, associate each virtual network that contains protected public IP addresses to the plan. Azure DDoS Protection is enabled at the virtual-network level, and eligible public IP resources in associated virtual networks receive the protection. Network security groups and Azure Firewall can provide complementary traffic filtering, but they do not provide the DDoS Protection plan benefits required in this scenario. See What is Azure DDoS Protection? and Manage Azure DDoS Protection.

QUESTION NO: 20

You have an Azure virtual network that contains two network virtual appliances (NVAs). The NVAs exchange routes with an on-premises router by using BGP.

You need to ensure that Azure virtual machines can dynamically learn routes advertised by the NVAs. The solution must minimize administrative effort.

What should you deploy?

A. Azure Route Server

B. Azure Traffic Manager

C. Azure NAT Gateway

D. Azure Front Door

ANSWER: A

Explanation:

Azure Route Server is the appropriate service because it provides managed BGP route exchange between network virtual appliances and Azure virtual network routing. Route Server establishes BGP peerings with the NVAs and can propagate routes learned from them to virtual machines in the virtual network. It also advertises relevant virtual network routes to the NVAs.

This avoids manually maintaining user-defined routes as on-premises prefixes or NVA routing information changes. Azure Route Server is deployed in a dedicated subnet named `RouteServerSubnet` and requires a minimum subnet size of `/27`. See Azure Route Server overview.

QUESTION NO: 21

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 Sql.EastUS

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

C. a deny rule that has a source of VirtualNetwork and a destination of 168.63.129.0/24

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

ANSWER: A D

Explanation:

An allow rule that has the IP address range of Vnet1 as the source and destination of Sql.EastUS permits outbound connections from the virtual machines to Azure SQL service endpoints in the East US region. `Sql.EastUS` is a regional service tag, so Azure maintains the underlying service IP prefixes and the NSG rule remains current as those prefixes change. The Vnet1 address range appropriately limits the rule's source to the virtual network's workloads.

A deny rule that has the IP address range of Vnet1 as the source and destination of Storage explicitly prevents those workloads from reaching Azure Storage endpoints. `Storage` is a service tag that represents Azure Storage address prefixes across Azure regions. An explicit outbound deny rule is required because NSGs include default outbound rules that otherwise permit broad traffic, including traffic to public Azure service endpoints. Assign the custom rules suitable priorities so that they are evaluated before any broader matching allow rule; NSG rules are processed in priority order, from the lowest numerical priority value upward. Together, these service-tag rules allow the required regional Azure SQL connectivity while blocking access to Azure Storage without manually managing Microsoft service IP ranges.

See Azure service tags overview and Network security groups overview.

QUESTION NO: 22

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:

2 is the minimum number of Azure DDoS Network Protection plans. A DDoS Network Protection plan is associated with virtual networks rather than being deployed separately for every public IP address or every protected resource. Therefore, the 20 internet-accessible resources in each virtual network do not require individual plans. One plan can protect the public IP resources in all virtual networks associated with that plan.

To reduce cost, a single plan should be shared by all eligible virtual networks that belong to subscriptions in the same Microsoft Entra tenant. Azure supports associating one DDoS Network Protection plan with multiple virtual networks, including virtual networks in different subscriptions, provided the subscriptions are in the same tenant. The subscription table represents subscriptions across two separate tenants, so a plan cannot be shared between those tenant boundaries. Deploying one plan for the virtual networks in each tenant protects all required resources while avoiding unnecessary additional plan charges.

Microsoft documents that DDoS Network Protection is enabled by associating a protection plan to virtual networks and that plans can be shared across subscriptions in the same tenant. See Manage Azure DDoS Protection and the Azure DDoS Protection overview.

QUESTION NO: 23

You have an Azure Front Door Premium profile named FD1. You add a custom domain named www.contoso.com to FD1.

Your organization requires that www.contoso.com use a certificate issued by a specified third-party certification authority. The certificate is stored in an Azure key vault named KV1.

You need to configure FD1 to use the certificate from KV1.

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

NOTE: Each correct selection is worth one point.

A. Enable a managed identity for FD1.

B. Assign the Key Vault Secrets User role for KV1 to the managed identity.

C. Configure the custom domain to use a customer-managed certificate from KV1.

D. Create a private endpoint for FD1 in KV1.

E. Enable a Front Door-managed certificate for the custom domain.

F. Assign the Reader role for KV1 to the managed identity.

ANSWER: A B C

Explanation:

Enable a **managed identity** for FD1 so that Azure Front Door has an Azure AD identity that can authenticate to Azure Key Vault. Grant that identity permission to read the certificate secret in KV1. With Azure RBAC, the **Key Vault Secrets User** role provides the required secret-read access for the Front Door identity.

You must then configure the custom domain HTTPS settings to use a **customer-managed certificate** and select the certificate in KV1. This allows Front Door to terminate TLS for the custom domain by using the required third-party CA-issued certificate. A Front Door-managed certificate cannot satisfy a requirement to use a certificate from a particular external certification authority.

See Configure HTTPS on an Azure Front Door custom domain and Manage Azure Front Door Standard/Premium certificates.

QUESTION NO: 24

You have an Azure App Service app named App1. App1 must access resources in an Azure virtual network named Vnet1 and an on-premises network connected to Vnet1 by using a VPN gateway.

You need to ensure that all outbound traffic from App1, including traffic to the on-premises network, is sent through Vnet1.

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

NOTE: Each correct selection is worth one point.

A. Configure regional virtual network integration for App1.

B. Enable Route All for App1.

C. Create a private endpoint for App1.

D. Enable a service endpoint on the integration subnet.

E. Configure an inbound NAT rule on the VPN gateway.

ANSWER: A B

Explanation:

Configure **regional virtual network integration** for App1 by using a dedicated integration subnet in Vnet1. Virtual network integration provides outbound connectivity from an App Service app to resources reachable through the integrated virtual network, including networks reachable through a VPN gateway.

You must also enable **Route All** for the virtual network integration. Route All causes all outbound traffic from the app to be routed into the integrated virtual network rather than only traffic destined for private address ranges. Routing and network security controls in Vnet1 can then direct the traffic to the VPN gateway and the on-premises network as required.

App Service virtual network integration is an outbound connectivity feature. It does not provide private inbound access to App1; a private endpoint is used when private inbound access is required. See Integrate your app with an Azure virtual network.

QUESTION NO: 25

N NO: 1

You need to configure the default route on Vnet2 and Vnet3. The solution must meet the virtual networking requirements.

What should you use to configure the default route?

A. route filters

B. BGP route exchange

C. a user-defined route assigned to GatewaySubnet in Vnet1

D. a user-defined route assigned to GatewaySubnet in Vnet2 and Vnet3

ANSWER: B

Explanation:

BGP route exchange is the appropriate mechanism for configuring and distributing a default route dynamically when the virtual network design uses a gateway-connected topology. BGP enables Azure VPN Gateway or ExpressRoute Gateway to exchange routes with its BGP peer, including the IPv4 default prefix, `0.0.0.0/0`, when that prefix is advertised by the connected network. Azure can then propagate the learned route through the applicable virtual-network gateway routing domain, allowing Vnet2 and Vnet3 to use the gateway path without maintaining separate static default-route entries.

This approach supports scalable routing: route changes can be advertised and withdrawn automatically through BGP instead of requiring manual changes to route tables as network prefixes or next hops evolve. Azure documents that routes learned

through BGP are included in a virtual network's effective routes when gateway route propagation is enabled. Default-route advertisement must be planned carefully because it directs traffic not matched by more-specific routes to the advertising gateway.

For Azure's route-selection behavior and BGP route propagation, see Virtual network traffic routing and About BGP with Azure VPN Gateway.

QUESTION NO: 26

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:

Azure Firewall Premium performs TLS inspection by decrypting outbound TLS sessions, inspecting the decrypted traffic according to firewall rules, and then re-encrypting it before sending it to the destination. To do this securely, the firewall requires an intermediate certificate authority certificate. **an Azure key vault** is the Azure service that stores this certificate and its private key. The firewall policy is configured to use the certificate held in the key vault for TLS inspection.

a managed identity, specifically a user-assigned managed identity, provides Azure Firewall with a secure identity for accessing the certificate secret in the key vault. The identity is assigned to the firewall and granted the required Key Vault permissions, allowing certificate retrieval without placing credentials or secret values in the firewall configuration. Client devices, including the Azure Virtual Desktop session hosts, must trust the corresponding root CA certificate so that inspected TLS connections continue to validate correctly. After these resources are configured, outbound traffic from the session-host subnet can be directed through Azure Firewall using routing, and the applicable Firewall Premium policy rules can enable TLS inspection.

See Azure Firewall Premium TLS inspection and Azure Firewall Premium certificates.

QUESTION NO: 27

You have an Azure application gateway for a web app named App1. The application gateway allows end-to-end encryption.

You configure the listener for HTTPS by uploading an enterprise-signed certificate.

You need to ensure that the application gateway can provide end-to-end encryption for App1.

What should you do?

A. Increase the Unhealthy threshold setting in the custom probe.

B. Enable the SSL profile to the listener.

C. Set Listener type to Multi site.

D. Upload the public key certificate to the HTTP settings.

ANSWER: D

Explanation:

Upload the public key certificate to the HTTP settings. End-to-end TLS with Azure Application Gateway requires encryption on both connection legs: client to the HTTPS listener, and Application Gateway to the backend web application. The certificate configured on the listener secures the frontend connection only. To establish a secure HTTPS connection to App1, the backend HTTP settings must use HTTPS and Application Gateway must be able to validate the certificate that the backend server presents during the TLS handshake.

For a backend certificate issued by a private or enterprise certificate authority, configure the trusted root certificate in the backend HTTP settings. This certificate establishes the trust chain Application Gateway uses when validating the backend server certificate. The certificate's subject name or SAN must also match the hostname Application Gateway uses for the backend connection. Once the backend certificate is trusted and HTTPS is selected in the HTTP settings, Application Gateway can decrypt the incoming HTTPS request as needed for gateway processing and then re-encrypt traffic to App1, completing end-to-end encryption. Microsoft describes this configuration in its Application Gateway TLS overview and end-to-end TLS configuration guidance.

QUESTION NO: 28 - (SIMULATION)

You are implementing the Virtual network requirements for Vnet6.

What is the minimum number of subnets and service endpoints you should create? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

ANSWER: Old Answer:

Answer:

Answer: 2, 4
New Answer:
Select 3 in the Subnets control and 2 in the Service endpoints control.Reasoning:
The image shows two numeric hotspot/dropdown controls in the Answer Area: one for Subnets and one for Service endpoints. The required virtual-network design uses three separately delegated or reserved network segments: a workload subnet, the required AzureFirewallSubnet, and the required GatewaySubnet. These Azure-managed gateway and firewall roles must each be placed in their own specifically named subnet and cannot share the workload subnet. The workload subnet must expose private, optimized access to the two required Azure PaaS service types by enabling the Microsoft.Storage and Microsoft.Sql service endpoints. Therefore, the correct selections are Subnets = 3 and Service endpoints = 2.

Explanation:

For Vnet6, select **3** for Subnets and **2** for Service endpoints. The minimum subnet count accounts for the separate network boundaries required by the design. A workload subnet is needed for the application resources. In addition, Azure Firewall requires its own dedicated subnet with the reserved name `AzureFirewallSubnet`, and a virtual network gateway requires its own dedicated subnet with the reserved name `GatewaySubnet`. These reserved subnets are separate from the workload subnet, so the three roles require three subnets in total.

The service-endpoint count is two because the workload subnet must enable endpoint connectivity for the two Azure service families required by the design: Azure Storage and Azure SQL Database. The endpoint service names are `Microsoft.Storage` and `Microsoft.Sql`. A service endpoint extends the VNet identity to the Azure service, allowing the corresponding service firewall to restrict access to selected virtual-network subnets while traffic remains on the Azure backbone. Endpoint enablement is configured on the subnet that contains the clients accessing those PaaS services.

Azure Firewall subnet requirements are documented by Microsoft at Deploy and configure Azure Firewall. Microsoft documents the dedicated gateway-subnet requirement at About VPN Gateway configuration settings. For the supported service endpoint service names and subnet configuration model, see Virtual Network service endpoints.

QUESTION NO: 29

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 C D

Explanation:

Four listed services require a dedicated subnet, so the stated requirement to select three is inconsistent with current Azure service requirements. A VPN gateway must be placed in a subnet named **GatewaySubnet**, which Azure reserves for gateway instances and gateway-related routing. Azure Bastion must be deployed in the dedicated **AzureBastionSubnet**; this subnet is reserved for the Bastion service. Microsoft Entra Domain Services (formerly Azure AD DS) also requires a dedicated subnet containing only the managed domain resources, allowing Azure to manage the domain controllers and associated networking safely. Azure Application Gateway v2 requires its own subnet for Application Gateway instances and service operations, including scaling and updates. These dedicated-subnet requirements ensure the respective Azure-managed services have the address space, routing control, and isolation they need. Microsoft documents the subnet requirements for VPN Gateway, Bastion, and Microsoft Entra Domain Services in its Virtual network integration guidance. Application Gateway subnet isolation and sizing requirements are documented in the Application Gateway infrastructure configuration guidance.

QUESTION NO: 30

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:

On-premises DNS servers must be configured with forwarders that point to the frontend IP address of FW1, and DNS proxy must be enabled on FW1. Together, these settings make Azure Firewall act as a DNS proxy endpoint reachable from the on-premises network over ExpressRoute. The on-premises DNS infrastructure can forward queries for `contoso.com` to the firewall's private/frontend address, rather than attempting to query Azure's platform DNS address directly.

With DNS proxy enabled, Azure Firewall accepts DNS requests on its private IP address and relays them to its configured upstream DNS service. When Azure-provided DNS is used as the upstream resolver, it can resolve records in Azure Private DNS zones that are linked to the firewall's virtual network. Because `contoso.com` is linked to Vnet1, the Azure DNS resolution path has visibility of the VM records registered in that private zone. This provides a controlled DNS-resolution bridge between the on-premises environment and Azure private DNS, without exposing Azure platform DNS directly to the on-premises network.

Microsoft documents Azure Firewall DNS proxy behavior and its use as a DNS forwarder at Azure Firewall DNS settings. For private DNS name resolution from on-premises networks, see Azure Private DNS zones.

QUESTION NO: 31

You need to connect Vnet2 and Vnet3. The solution must meet the virtual networking requirements and the business requirements. Which two actions should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. On the peerings from Vnet2 and Vnet3, select Use remote gateways.

B. On the peering from Vnet1, select Allow forwarded traffic.

C. On the peering from Vnet1, select Use remote gateways.

D. On the peering from Vnet1, select Allow gateway transit.

E. On the peerings from Vnet2 and Vnet3, select Allow gateway transit.

ANSWER: A D

Explanation:

"On the peerings from Vnet2 and Vnet3, select Use remote gateways" and "On the peering from Vnet1, select Allow gateway transit" configure gateway transit through the shared gateway in Vnet1. Gateway transit lets a hub virtual network make its virtual network gateway available to networks that are connected to it through virtual network peering. Enabling *Allow gateway transit* on Vnet1's peering advertises the hub gateway for use by the connected spoke networks. Enabling *Use remote gateways* on the Vnet2 and Vnet3 peerings causes each spoke to consume that advertised gateway rather than requiring its own gateway. This combination supports a centralized hub-and-spoke design, including shared connectivity to networks reachable through the hub gateway, while avoiding the cost and operational overhead of separate gateways in each spoke. Azure requires the gateway-transit setting on the gateway-owning virtual network and the remote-gateway setting on the virtual networks that use it. A virtual network configured to use a remote gateway cannot also have its own gateway. See Azure virtual network peering gateway connectivity and Configure gateway transit for virtual network peering.

QUESTION NO: 32

You have an Azure virtual network named Vnet1 that is linked to a private DNS zone named contoso.com.

You deploy 20 virtual machines to Vnet1. You need the virtual machines to automatically register their host names and private IP addresses in the contoso.com zone.

What should you configure?

A. Enable auto-registration on the Vnet1 link.

B. Create an Azure DNS Private Resolver outbound endpoint.

C. Enable DNS proxy on Azure Firewall.

D. Create a service endpoint for Azure DNS.

ANSWER: A

Explanation:

You should **enable auto-registration on the virtual network link** for the private DNS zone. When registration is enabled for a virtual network link, Azure automatically creates and maintains DNS A records for virtual machines deployed in the linked virtual network.

Only one private DNS zone link with auto-registration enabled can exist for a given virtual network. The same virtual network can be linked to additional private DNS zones for resolution, but those links cannot also enable auto-registration. See Azure Private DNS virtual network links.

QUESTION NO: 33 - (DRAG DROP)

You need to prepare Vnet1 for the deployment of an ExpressRoute gateway. The solution must meet the hybrid connectivity requirements and the business requirements.

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

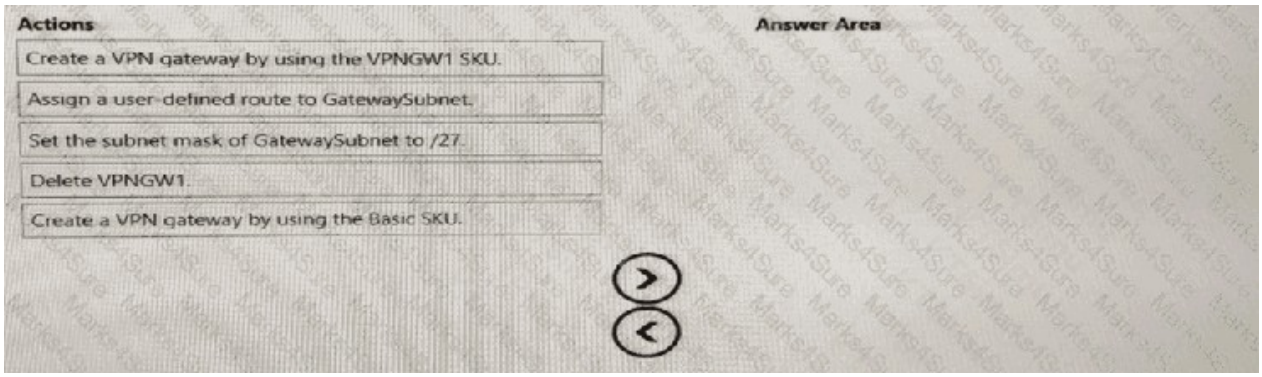

ANSWER:

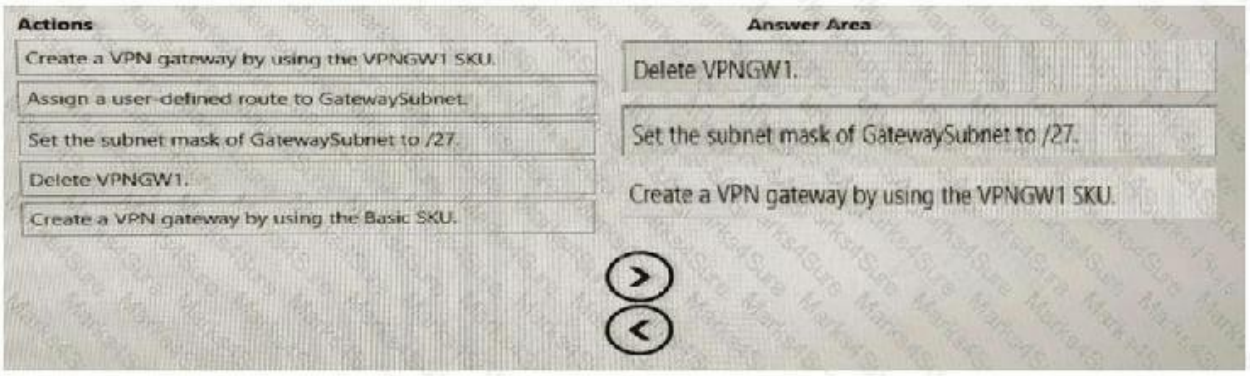

Explanation:

An ExpressRoute virtual network gateway is deployed into the dedicated subnet named **GatewaySubnet**. Before the subnet can be changed, the existing gateway, VPNGW1, must be deleted. Azure does not allow changing the address range of GatewaySubnet while a virtual network gateway is using that subnet. Deleting VPNGW1 first releases the subnet and allows its configuration to be updated safely.

Next, GatewaySubnet should be assigned a **/27** subnet mask. Gateway gateways consume multiple IP addresses for their gateway instances and supporting infrastructure. A /27 provides sufficient address capacity and is the recommended subnet size for gateways, particularly where the solution must accommodate hybrid connectivity requirements and avoid limitations as the gateway is deployed or maintained. The subnet must retain the exact reserved name, GatewaySubnet, because Azure identifies it by name during gateway deployment.

After GatewaySubnet has been resized, create the new VPN gateway using the **VpnGw1** SKU. This SKU provides the supported gateway capabilities needed for an ExpressRoute gateway deployment, unlike the legacy Basic gateway SKU. Creating the gateway last ensures that Azure deploys it into the correctly sized GatewaySubnet and that the gateway meets the required hybrid connectivity design.

Microsoft documents the GatewaySubnet requirements and recommends allocating enough IP addresses for gateway instances and future needs in its Gateway subnet documentation. The supported gateway SKU capabilities, including ExpressRoute support, are described in the About ExpressRoute virtual network gateways documentation.

QUESTION NO: 34 - (HOTSPOT)

You ate configuring the DNS forwarding luleset for DNSR1

You need to configure the destination IP address for azure.proseware.com and for corp.proseware.com. The solution must meet the general requirements.

Which IP addiesses should you configure for each namespace? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

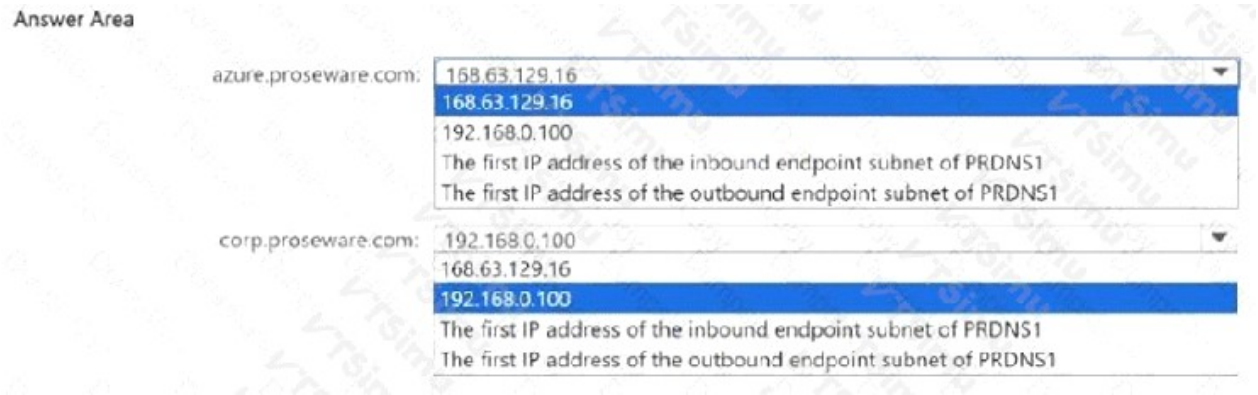

ANSWER:

Explanation:

The forwarding rule for **azure.proseware.com** should use **168.63.129.16**. This address is the Azure platform virtual IP address that provides DNS services to resources within Azure virtual networks. It is the Azure-provided DNS service address used to resolve Azure-hosted names, so it is the appropriate destination for queries in the Azure namespace. Azure documents this virtual IP as part of the platform services available in every virtual network; see What is IP address 168.63.129.16?.

The forwarding rule for **corp.proseware.com** should use **192.168.0.100**. This is the corporate DNS server address, so forwarding queries for the corporate namespace to that server ensures they are resolved by the DNS infrastructure that hosts and manages corporate DNS records. In an Azure DNS Private Resolver forwarding ruleset, each rule matches a DNS suffix and sends matching queries to the configured target DNS server IP address. This lets Azure workloads resolve names that belong to a connected corporate environment while retaining Azure DNS resolution for the Azure namespace. The forwarding-rule behavior and use of target DNS server addresses are described in Azure DNS Private Resolver forwarding rulesets and rules.

Therefore, configure **168.63.129.16** for **azure.proseware.com** and **192.168.0.100** for **corp.proseware.com**.

QUESTION NO: 35

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:

NSG flow logs should be implemented first because Traffic Analytics depends on the network-flow telemetry that these logs collect. NSG flow logs record IP traffic flowing through network security groups associated with Azure virtual-machine network interfaces or subnets. The records include information such as source and destination IP addresses, ports, protocol, whether the flow was allowed or denied, and traffic-volume details. Traffic Analytics ingests and aggregates this flow-log data in a Log Analytics workspace to present visualizations and insights about traffic patterns, communicating hosts, ports, applications, and potential security issues.

In practice, configure NSG flow logs for the relevant NSGs, associate the configuration with a storage account and Log Analytics workspace as required, and enable Traffic Analytics for the flow logs. After flow data is collected and processed, Traffic Analytics can provide the application-usage visibility requested. Microsoft documents that Traffic Analytics is a cloud-based solution that provides insights by analyzing NSG flow logs. See Traffic Analytics documentation and NSG flow logs overview.

QUESTION NO: 36

You plan to deploy Azure Route Server in a virtual network named Vnet1. Two network virtual appliances will exchange routes with Route Server by using BGP.

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 dedicated subnet named RouteServerSubnet that is at least /27.

B. Create a BGP peering between Azure Route Server and each NVA.

C. Associate a route table to GatewaySubnet.

D. Create a local network gateway for each NVA.

E. Enable service endpoints on the NVA subnet.

ANSWER: A B

Explanation:

Azure Route Server must be deployed to a dedicated subnet named **RouteServerSubnet**. The subnet must be at least `/27`, and it cannot contain other resources. This provides the required address capacity for the Route Server service.

You must also create a **BGP peering** between Azure Route Server and each NVA. The peering configuration specifies the NVA IP address and autonomous system number. After the BGP sessions are established, Route Server can exchange routes dynamically with the NVAs. See Configure Azure Route Server and Azure Route Server overview.

QUESTION NO: 37

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 configure the firewall on storage1 to only accept connections from Vnet1. Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. An Azure Storage firewall is an inbound access control applied at the individual storage account. Allowing traffic from Vnet1 on storage1 can permit VM1 to reach storage1, provided the appropriate network path is configured, such as a service endpoint from Subnet1 or a private endpoint. However, the stated requirement also imposes an outbound-access restriction: VM1 must not be able to access any storage account other than storage1. Configuring the network rules of storage1 does not govern destinations that VM1 can initiate connections to, including storage2, storage3, or any other Azure Storage public endpoint that remains reachable.

Meeting the complete requirement requires an egress-control design for VM1, such as routing its traffic through Azure Firewall or another network virtual appliance and allowing only the intended storage destination. A private endpoint for storage1, together with private DNS and appropriately restricted outbound paths, can also be part of this design. Azure Storage network rules are evaluated by each storage account to determine whether that account accepts a request; they are not a mechanism for constraining all storage destinations available to a virtual machine. See Azure Storage network security and Azure Firewall FQDN tags.

QUESTION NO: 38

You plan to deploy Azure Firewall Premium in a virtual network named Vnet1.

All traffic from workload subnets must be routed through the firewall. The firewall must use forced tunneling to send Internet-bound traffic to an on-premises security appliance by using a VPN gateway.

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

NOTE: Each correct selection is worth one point.

A. An AzureFirewallManagementSubnet subnet

B. A public IP address for the firewall management configuration

C. An AzureBastionSubnet subnet

D. An Application Gateway v2

E. An Azure NAT Gateway

ANSWER: A B

Explanation:

AzureFirewallManagementSubnet is required when Azure Firewall is deployed in forced-tunneling mode. The firewall uses this dedicated subnet for management traffic. The subnet must be named exactly `AzureFirewallManagementSubnet` and must have a minimum size of `/26`.

A **public IP address for the management configuration** is also required. Azure Firewall uses the management public IP address for platform management traffic even when data-plane Internet-bound traffic is forced through a virtual network gateway to an on-premises appliance. The standard public IP address used for the firewall data plane is separate from this management configuration. See Azure Firewall forced tunneling and Deploy and configure Azure Firewall.

QUESTION NO: 39

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: C

Explanation:

a listener, a backend pool, and a rule is required to publish App2 through the existing Azure Application Gateway. Application Gateway processes a request by first accepting it through a listener, then evaluating a request-routing rule, which sends the request to the backend pool configured for the target application. For a second web app accessed by its own hostname, a new multi-site listener provides the hostname match for App2 while allowing Gateway1 to continue serving App1 on the same frontend IP address and port.

App2 must also be represented by a backend pool. The backend pool identifies the App Service endpoint to which Application Gateway forwards requests. Finally, a request-routing rule associates the new listener with the backend pool and the applicable backend HTTP settings. This creates the complete listener-to-rule-to-backend configuration path needed for Application Gateway to route requests specifically to App2, without altering the existing App1 publishing configuration.

Microsoft documents that listeners accept incoming traffic and that request-routing rules direct listener traffic to backend pools. Multi-site listeners support hosting multiple applications by using host names on a shared Application Gateway frontend. See Application Gateway configuration overview and multiple-site hosting with Application Gateway.

QUESTION NO: 40

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 required because it provides the vendor-specific SD-WAN functionality needed for Azure to participate in an existing third-party SD-WAN fabric. An NVA is a virtual appliance supplied by, or configured for, the SD-WAN vendor and can implement the SD-WAN overlay, control-plane integration, routing behavior, policy enforcement, and orchestration required to make the Azure deployment function as a node of that vendor's solution. Azure Virtual WAN supplies the managed connectivity and routing foundation, including a virtual hub and connectivity for virtual networks, branches, and gateways. However, the third-party SD-WAN capabilities themselves are implemented by the vendor appliance rather than by a Virtual WAN gateway alone. The NVA can be connected into the Virtual WAN architecture so traffic between the Azure virtual networks and on-premises SD-WAN sites follows the organization's SD-WAN design and policies. Microsoft documents Network Virtual Appliances as supported components in Virtual WAN architectures and describes Virtual WAN as an Azure networking service that integrates branches and virtual networks through virtual hubs. See About Azure Virtual WAN and About Network Virtual Appliance integration in Azure Virtual WAN.

QUESTION NO: 41

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. Modify the DNS server settings of Vnet1.

B. For FW1, configure custom DNS server.

C. For FW1, enable DNS proxy.

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

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

ANSWER: C D

Explanation:

Enabling DNS proxy on FW1 makes the firewall available as a DNS proxy endpoint. It receives DNS queries on its private IP address and forwards them to its configured upstream resolver. When Azure Firewall uses Azure-provided DNS resolution, it

can resolve records hosted in Azure Private DNS zones that are linked to its virtual network, including the *contoso.com* zone linked to Vnet1. This provides the Azure-side DNS resolution component needed for private zone records.

Configuring forwarders on the on-premises DNS servers that point to the frontend IP address of FW1 provides the required query path from the datacenter over ExpressRoute. In practice, the forwarder should be configured as a conditional forwarder for *contoso.com* and use the firewall's reachable private IP address. Queries for names in that zone are then sent to FW1, which proxies the queries to Azure DNS and returns the private DNS responses to the on-premises DNS servers. This approach keeps the Azure-provided DNS address internal to Azure while exposing a supported, reachable DNS proxy endpoint to the on-premises network.

For implementation details, see Azure Firewall DNS settings and Azure Private Endpoint DNS integration.

QUESTION NO: 42

You plan to use Azure Virtual WAN to connect branch offices to virtual networks in multiple Azure regions.

You need to provide managed transit connectivity between the branch offices and the virtual networks.

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

NOTE: Each correct selection is worth one point.

A. an Azure Virtual WAN

B. an Azure Virtual Hub

C. an application security group

D. an Azure private DNS zone

E. an Azure Front Door profile

ANSWER: A B

Explanation:

An **Azure Virtual WAN** is the top-level networking resource that provides the management and orchestration plane for a Virtual WAN deployment. It can contain multiple virtual hubs across Azure regions.

An **Azure Virtual Hub** provides the regional transit hub where virtual network connections, VPN sites, ExpressRoute circuits, and supported security services can connect. Virtual WAN routes traffic through the virtual hubs over the Microsoft global network, which provides managed branch-to-VNet and branch-to-branch connectivity. See What is Azure Virtual WAN? and About Virtual WAN hubs.

QUESTION NO: 43 - (SIMULATION)

Task 2

You need to ensure that you can deploy Azure virtual machines to the France Central Azure region. The solution must ensure that virtual machines in the France Central region are in a network segment that has an IP address range of 10.5.1.0.

ANSWER: See the explanation for the answer

Explanation:

Create a virtual network in the France Central region and create a subnet whose prefix is `10.5.1.0/24`.

In the Azure portal, go to **Virtual networks** and select **Create**.
Select the required subscription and resource group.
Set **Region** to **France Central**.
Configure a VNet address space that contains the required subnet, for example `10.5.0.0/16`.

Add a subnet (any appropriate name) with subnet address range `10.5.1.0/24`.
Create the VNet.

When creating VMs in France Central, select this VNet and the `10.5.1.0/24` subnet on the **Networking** tab. Azure requires CIDR notation; therefore, `10.5.1.0/24`, rather than only `10.5.1.0`, is the required subnet range.

QUESTION NO: 44

You have an Azure Application Gateway v2 named AppGW1. AppGW1 has a backend pool that contains two web servers.

Users report that AppGW1 intermittently returns HTTP 502 errors.

You need to identify the reason that Application Gateway considers the backend servers unhealthy.

What should you review?

A. Backend health

B. Effective routes

C. NSG flow logs

D. Azure Advisor

ANSWER: A

Explanation:

You should review the **backend health** status of AppGW1. Application Gateway uses health probes to determine whether backend targets can accept requests. Backend health identifies each backend member as healthy, unhealthy, or unknown and provides diagnostic details, such as probe timeout, DNS resolution failure, TLS validation failure, or an unexpected HTTP response status code.

An HTTP 502 response from Application Gateway commonly indicates that no healthy backend server is available for the matching routing rule. Reviewing backend health identifies the failed probe condition so that the backend HTTP settings, probe configuration, certificates, DNS, or server application can be corrected. See Application Gateway backend health.

QUESTION NO: 45

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

* An Azure App Service app named App1

* An Azure DNS zone named contoso.com

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

* A virtual network named Vnet1

You create a private endpoint for App1. The record for the endpoint is registered automatically in Azure DNS.

You need to provide a developer with the name that is registered in Azure DNS for the private endpoint.

What should you provide?

A. app1.privatelink.azurewebsites.net

B. app1.contoso.com

C. app1.contoso.onmicrosoft.com

D. app1.private.contoso.com

ANSWER: A

Explanation:

app1.privatelink.azurewebsites.net is the private DNS name registered for an Azure App Service app named App1 when it is exposed through an Azure Private Endpoint. Azure Private Link uses service-specific private DNS zones. For Microsoft.Web sites, the designated private DNS zone is **privatelink.azurewebsites.net**. The private endpoint DNS record uses the App Service app name as the record name, producing the fully qualified name **app1.privatelink.azurewebsites.net**.

This private DNS record is an A record that resolves to the private IP address assigned to the private endpoint. When the appropriate private DNS zone is linked to a virtual network, workloads in that network can resolve the App Service endpoint privately. App Service's normal public hostname resolution is designed to accommodate this Private Link DNS naming pattern, allowing private connectivity without exposing traffic through the public endpoint.

The relevant DNS suffix is determined by the Azure resource type, rather than by the Azure AD tenant name, a public DNS zone in the subscription, or an arbitrary existing private DNS zone. Microsoft lists **privatelink.azurewebsites.net** as the private DNS zone for Azure App Service private endpoints. See Azure Private Endpoint private DNS zone values and Use private endpoints for Azure App Service apps.

QUESTION NO: 46

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 Sql.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 168.63.129.0/24

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

ANSWER: A D

Explanation:

An allow rule with the Vnet1 address range as the source and `Sql.EastUS` as the destination permits outbound traffic from the virtual machines to Azure SQL service endpoints identified for the East US region. Azure service tags represent Microsoft-managed groups of IP prefixes, so the rule does not require manually maintaining the changing public address ranges used by the Azure SQL service. The regional `Sql.EastUS` tag scopes that permitted SQL connectivity to the specified Azure region.

A deny rule with the Vnet1 address range as the source and `Storage` as the destination blocks outbound connectivity to Azure Storage service endpoints. The `Storage` service tag represents Azure Storage IP prefixes and is appropriate when access to storage resources must be prevented. Configure this deny rule with a higher priority than any broad outbound allow rule that could otherwise match the traffic. NSG rules are evaluated by priority, with lower numeric values processed first; a matching deny rule therefore prevents the storage connection while the SQL-specific allow rule supports the required East US SQL traffic. See Azure service tags overview and Network security groups overview.

QUESTION NO: 47

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 B D E

Explanation:

Azure Bastion requires a dedicated subnet named **AzureBastionSubnet**. This subnet is reserved for the Bastion service and must be sized to support the selected Bastion SKU and any required scale units. A VPN gateway similarly requires a dedicated subnet named **GatewaySubnet**, where Azure deploys and manages the gateway instances. Address planning must reserve this subnet separately from workloads because no other resources can be deployed there.

Azure Application Gateway v2 requires a dedicated subnet. Application Gateway instances, their private frontend configuration, and scaling capacity are placed in that subnet, so it must not be shared with other resource types. Azure Active Directory Domain Services also requires deployment into a dedicated subnet: the managed domain controllers use IP addresses from that subnet, and the subnet must not contain other workloads or gateway resources. Therefore, although the question requests three selections, four listed services have dedicated-subnet requirements. Reserve address space for each of these platform-managed services when designing the virtual network. See Azure Bastion configuration settings and Microsoft Entra Domain Services network considerations.

QUESTION NO: 48

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:

Use a custom WAF rule to evaluate the incoming request header and perform the required redirect when that header contains `string2`. Azure Front Door WAF custom rules support matching against request headers and can apply a redirect action, including a destination URL. The rule should be configured with a request-header match condition using the appropriate header name and a *Contains* match operator, then configured to redirect to `https://www.contoso.com/redirect2`. Creating the custom rule defines this additional header-based behavior, and adding the custom rule to Policy1 ensures that it is processed alongside the existing rule for `string1`. The rule priority should be selected so its processing order is intentional if requests can satisfy more than one rule. Creating an association binds the WAF policy to the applicable Azure Front Door frontend, allowing Policy1 and its custom rules to inspect requests received by Frontend1. In this scenario, the association should continue to target Frontend1 after the policy update. Microsoft documents request-header matching and redirect actions for Front Door WAF custom rules at Microsoft Learn: Azure Front Door WAF custom rules. Policy association and configuration are described in Microsoft Learn: Create an Azure Front Door WAF policy.

QUESTION NO: 49

You have an Azure subscription that contains a virtual network name Vnet1. Vnet1 contains a virtual machine named VM1 and an Azure firewall named FW1.

You have an Azure Firewall Policy named FP1 that is associated to FW1.

You need to ensure that RDP requests to the public IP address of FW1 route to VM1. What should you configure on FP1?

A. an application rule

B. a network rule

C. a DNAT rule

ANSWER: C

Explanation:

A DNAT rule is the required Azure Firewall Policy configuration because the requirement is to publish an internal resource through the firewall's public IP address. A destination network address translation rule listens for traffic addressed to a specified Azure Firewall public IP address and port, then translates that destination to the private IP address and port of the target workload. For this scenario, the rule would use TCP protocol, destination port 3389, and VM1's private IP address as the translated address. The translated port would normally also be 3389. DNAT rules are configured in a NAT rule collection within the firewall policy. Azure Firewall processes NAT rules before network rules, allowing the incoming RDP session to be translated and forwarded to the internal virtual machine. This approach permits controlled inbound administrative access without assigning a public IP address directly to VM1. Microsoft's Azure Firewall DNAT tutorial describes creating a NAT rule collection with the firewall public IP, an inbound port, and the translated private address and port. For implementation guidance, see Tutorial: Deploy and configure Azure Firewall using the Azure portal and Azure Firewall Policy rule sets.

QUESTION NO: 50

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 Sql.EastUS

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

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:

An allow rule that uses Vnet1's IP address range as its source and `Sql.EastUS` as its destination permits outbound traffic from the virtual machines to Azure SQL service IP addresses published for the East US region. Azure service tags are Microsoft-managed groups of IP prefixes, so using the regional SQL service tag avoids maintaining changing service IP ranges manually. The source can be the Vnet1 address range, which scopes the rule to the workloads in that virtual network.

A deny rule that uses Vnet1's IP address range as its source and `Storage` as its destination blocks outbound traffic from those virtual machines to the Azure Storage service tag. The `Storage` tag represents Azure Storage service address prefixes and is maintained by Microsoft as those prefixes change. The deny rule must have a priority that causes it to be evaluated before any broader outbound allow rule that could otherwise permit the traffic. Together, the regional SQL allow rule and Storage deny rule provide the requested service-tag-based control for the specified SQL and Storage destinations.

For service tag syntax and the supported regional tags, see Azure service tags overview. NSG outbound rule processing and priority behavior are documented in Network security groups.

QUESTION NO: 51

You have an Azure subscription that contains a resource group named RG1 and a virtual network named VNet1 You need to deploy Azure Firewall to RG1. The solution must minimize administrative effort What should you do first?

A. Create a secured virtual hub named AzureFirewallHub.

B. Create a new resource group named AzureFirewallResourceGroup.

C. Create a new virtual network named AzureFirewallNetwork.

D. On VNet1, create a virtual subnet named AzureFirewallSubnet.

ANSWER: D

Explanation:

On VNet1, create a virtual subnet named AzureFirewallSubnet. Azure Firewall is deployed into a customer-managed virtual network only when that virtual network contains a dedicated subnet with the exact required name, **AzureFirewallSubnet**. This subnet is reserved exclusively for the firewall and must have sufficient address space; Microsoft recommends a **/26** subnet or larger to support firewall infrastructure and future scaling. Because VNet1 already exists and the firewall is to be deployed in RG1, preparing the required subnet in VNet1 is the prerequisite that enables deployment without creating and managing separate networking resources. During deployment, Azure Firewall uses this subnet for its internal IP configuration, while a public IP resource is associated for internet-facing connectivity as needed. The subnet must be created before the firewall resource can be deployed into the virtual network. This directly satisfies the requirement to minimize administrative effort by using the existing resource group and virtual network rather than introducing a new virtual network or a Virtual WAN secured hub. For implementation requirements and deployment steps, see Deploy and configure Azure Firewall and Tutorial: Deploy and configure Azure Firewall using the Azure portal.

QUESTION NO: 52

You have an Azure subscription that contains a user named Admin1 and a resource group named RG1.

RG1 contains an Azure Network Watcher instance named NW1.

You need to ensure that Admin1 can place a lock on NW1. The solution must use the principle of least privilege.

Which role should you assign to Admin1?

A. User Access Administrator

B. Network Contributor

C. Resource Policy Contributor

D. Monitoring Contributor

ANSWER: C

Explanation:

Resource Policy Contributor is the least-privileged built-in Azure role that enables Admin1 to place a management lock on NW1. Azure resource locks are managed through Azure Resource Manager authorization operations, specifically the `Microsoft.Authorization/locks/*` permissions. The Resource Policy Contributor role includes permissions to create, update, and delete locks, allowing Admin1 to apply a *CanNotDelete* or *ReadOnly* lock at the Network Watcher resource scope. Assigning the role directly on NW1 limits its effective scope to that individual resource, which further follows least-privilege design.

Management locks protect resources from accidental modification or deletion, including changes made through the Azure portal, Azure CLI, Azure PowerShell, REST APIs, or SDKs. A lock applied directly to NW1 affects that resource without unnecessarily granting Admin1 permissions to administer other networking resources in RG1 or elsewhere in the subscription. Microsoft documents that creating or deleting resource locks requires the relevant `Microsoft.Authorization/locks/*` authorization permission, and Resource Policy Contributor is a built-in role intended for policy and lock management. See Lock your Azure resources to protect your infrastructure and Resource Policy Contributor built-in role.