PassQueen
a.k.a dumpsqueen.com

Administering Windows Server Hybrid Core Infrastructure

Microsoft AZ-800

Version Demo

Total Demo Questions: 33

Total Premium Questions: 332

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Deploy and manage Active Directory Domain Services (AD DS) in on-premises and cloud environments	113
Topic 2, Manage Windows Servers and workloads in a hybrid environment	38
Topic 3, Manage virtual machines and containers	51
Topic 4, Implement and manage an on-premises and hybrid networking infrastructure	66
Topic 5, Manage storage and file services	49
Topic 6, Case Study 1	5
Topic 7, Case Study 2	4
Topic 8, Case Study 3	2
Topic 9, Case Study 4	4
Total	332

QUESTION NO: 1

You have an Azure subscription that contains a Recovery Services vault named Vault1.

You have an on-premises server named Server1 that runs Windows Server. You need to back up files and folders on Server1 directly to Vault1.

Which two actions should you perform on Server1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Install the Microsoft Azure Recovery Services agent.

B. Register Server1 with Vault1 by using vault credentials.

C. Install the Azure VM backup extension.

D. Install the Azure Connected Machine agent.

E. Create a recovery plan in Vault1.

ANSWER: A B

Explanation:

To back up files and folders from an on-premises Windows Server directly to a Recovery Services vault, install the **Microsoft Azure Recovery Services (MARS) agent** on the server. The MARS agent supports file and folder backup to Azure Backup without requiring System Center Data Protection Manager or Azure Backup Server.

You must also register the MARS agent with Vault1 by using vault credentials downloaded from the Recovery Services vault. Registration establishes the trust relationship between Server1 and the vault and allows you to configure a backup schedule and retention policy. The Azure VM backup extension is only for Azure virtual machines. For more information, see Back up Windows Server files and folders to Azure.

QUESTION NO: 2

You have an on-premises Active Directory Domain Servic es (AD DS) domain that syncs with an Azure Active Directory (Azure AD) tenant Group writeback is enabled in Azure AD Connect.

The AD DS domain contains a server named Server1 Server 1 contains a shared folder named share1.

You have an Azure Storage account named storage2 that uses Azure AD-based access control. The storage2 account contains a share named shared

You need to create a security group that meets the following requirements:

• Can contain users from the AD DS domain

• Can be used to authorize user access to share 1 and share2

What should you do?

A. in the AD DS domain, create a universal security group

B. in the Azure AD tenant create a security group that has assigned membership

C. in the Azure AD Tenant create a security group that has dynamic me mbership.

D. in the Azure AD tenant create a Microsoft 365 group

ANSWER: B

Explanation:

In the Azure AD tenant, create a security group that has assigned membership. This provides a single cloud-managed security principal that can be used when assigning Azure role-based access control permissions for the Azure file share. Assigned membership is appropriate because administrators explicitly add the required users, including users synchronized from the on-premises AD DS domain, to the group.

With Azure AD Connect group writeback enabled, the cloud-created security group is written back to AD DS as a universal security group. The written-back group can then be selected when configuring share and NTFS permissions for the shared folder on Server1. As a result, the same group membership is used for authorization in both environments: Azure Storage through Azure AD-based access control and the on-premises file server through AD DS permissions. This avoids maintaining separate groups and independently reconciling membership for the cloud and on-premises resources.

Microsoft documents group writeback as a mechanism for making cloud groups available in on-premises Active Directory, including for access to on-premises resources. See Microsoft Entra Connect group writeback. Azure Files also supports assigning share-level permissions through Azure RBAC to Microsoft Entra users and groups, as described in Assign share-level permissions for Azure Files.

QUESTION NO: 3

You have an Azure subscription that contains a virtual machine named VM1 as shown in the following exhibit.

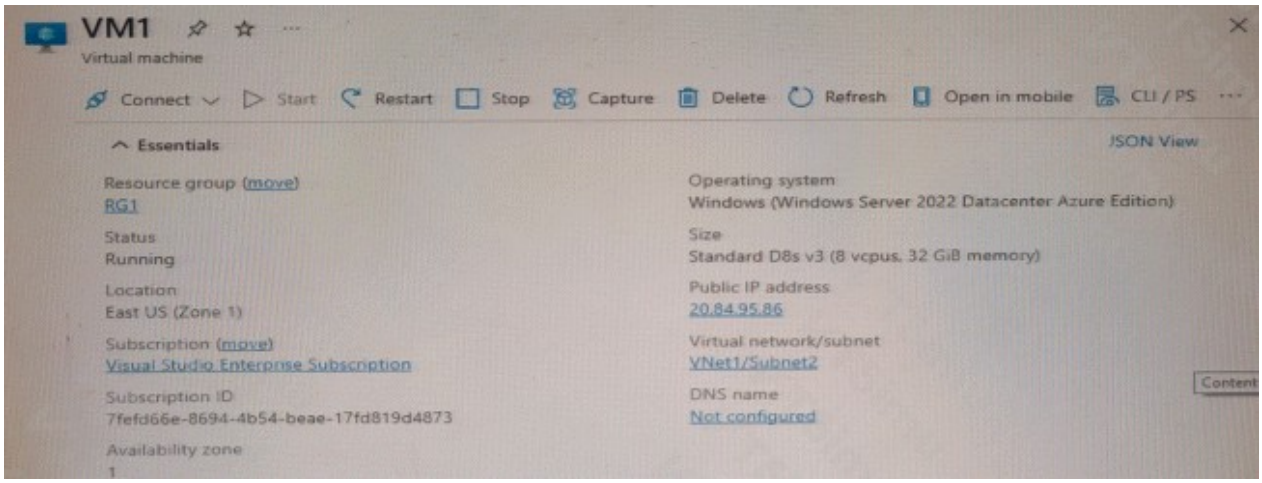

The subscription has the disks shown in the following table.

Name	Resource group	Location	Availability zone
Disk1	RG1	East US	None
Disk2	RG2	East US	1
Disk3	RG1	Central US	None
Disk4	RG1	Central US	1

Which disks can you attach as data disks to VM1?

A. Disk2 only

B. Disk4 only

C. Disk1 and Disk2 only

D. Disk2 and Disk4 only

E. Disk1, Disk3, and Disk4 only

F. Disk1, Disk2, Disk3. and Disk4

ANSWER: C

Explanation:

Disk1 and Disk2 only is correct. Azure data disks must be attachable to the target virtual machine's deployment context. In particular, a managed disk must be in the same Azure region as the VM, must not already be attached to another VM, and must meet any applicable availability-zone compatibility requirements. The disk configuration shown for Disk1 and Disk2 satisfies the requirements for attachment to VM1, so either disk can be added as a data disk.

When a disk is attached, Azure exposes it to the guest operating system as an additional disk. The administrator must then initialize, partition, format, and mount the disk in the operating system before it can store data. Azure supports attaching managed disks as data disks through the portal, Azure CLI, PowerShell, ARM templates, or other supported management tools. The VM size also determines the maximum number of data disks that can be attached, and the selected disks are within the applicable attachment constraints shown in the exhibits.

For the platform requirements and procedures for attaching a managed data disk, see Attach a managed data disk to a Windows VM and Azure managed disks overview.

QUESTION NO: 4

You have on-premises Windows devices.

You have an A zure subscription that contains a virtual network named VNet1.

You need to create a Site-to-Site (S2S) VPN between the on-premises network and VNet1.

Which three resources should you create? Each correct answer present part of the solution.

NOTE; Each corr ect selection is worth one point.

A. a network security group (NSG)

B. an Azure NAT gateway

C. an Azure Route Server

D. a VPN connection

E. a virtual network gateway

F. a local network gateway

ANSWER: D E F

Explanation:

A Site-to-Site VPN requires a VPN connection, a virtual network gateway, and a local network gateway. The virtual network gateway is the Azure-managed gateway resource deployed to the gateway subnet in VNet1. It provides the VPN endpoint on the Azure side and processes encrypted IPsec/IKE traffic between Azure and the on-premises network. The local network gateway represents the on-premises VPN device and on-premises network in Azure. Its configuration includes the public IP address of the on-premises VPN device and the address prefixes that are reachable in the local network. A VPN connection then links the virtual network gateway to the local network gateway and contains the shared key and IPsec/IKE configuration used to establish the encrypted tunnel. Together, these resources define both tunnel endpoints and the connection between them. The on-premises side must also have a compatible VPN device configured with corresponding settings, but that device is not an Azure resource created in the subscription. Microsoft describes this resource model in its VPN Gateway overview and provides the deployment workflow in Create a Site-to-Site connection using the Azure portal.

QUESTION NO: 5 - (DRAG DROP)

DRAG DROP

You deploy a new Active Directory Domain Services (AD DS) forest named contoso.com. The domain contains three domain controllers named DC1, DC2, and DC3.

You rename Default-First-Site-Name as Site1.

You plan to ship DC1, DC2, and DC3 to datacenters in different locations.

You need to configure replication between DC1, DC2, and DC3 to meet the following requirements:

- Each domain controller must reside in its own Active Directory site.
- The replication schedule between each site must be controlled independently.
- Interruptions to replication must be minimized.

Which three actions should you perform in sequence in the Active Directory Sites and Services console? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

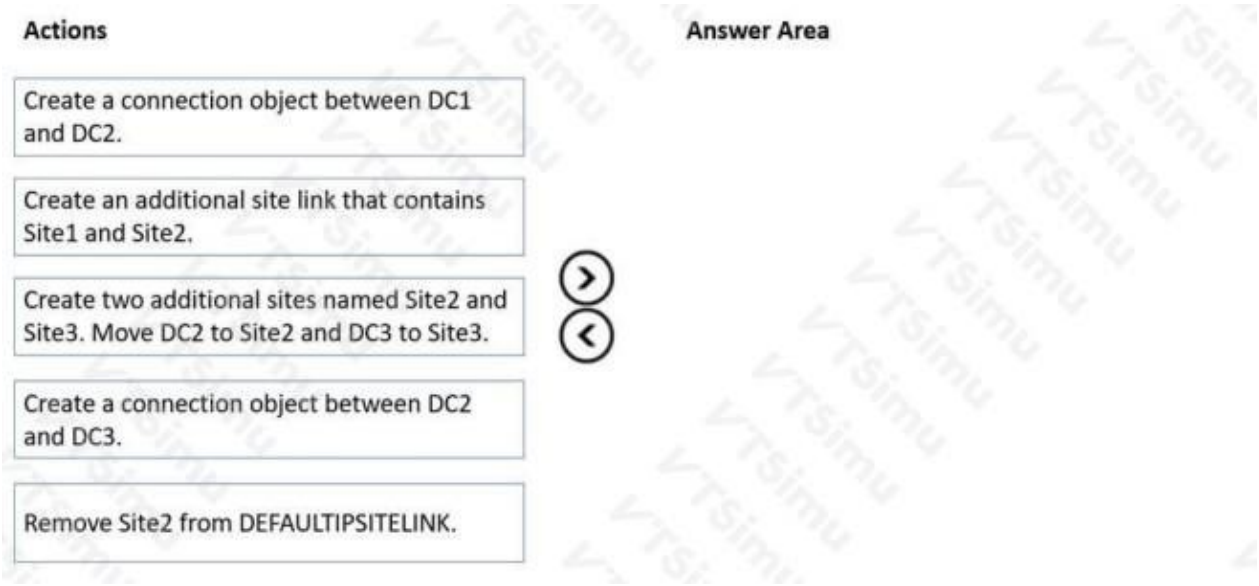

ANSWER:

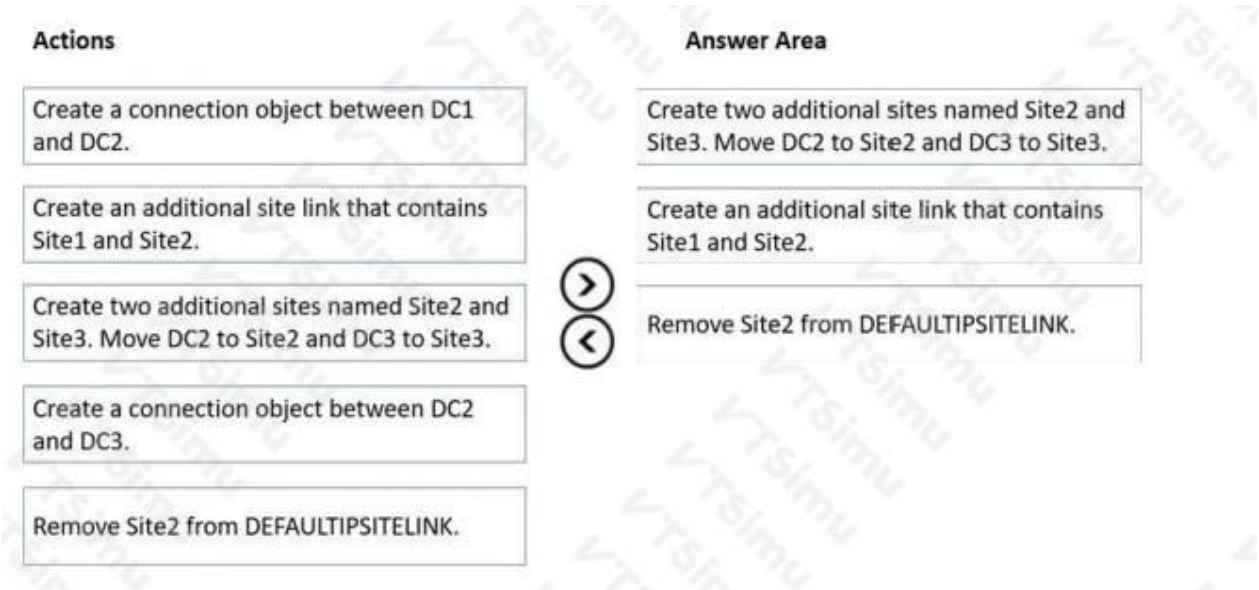

Explanation:

Begin by creating Site2 and Site3, then move DC2 to Site2 and DC3 to Site3. DC1 remains in Site1, so this establishes the required arrangement of one domain controller per Active Directory site. Site objects represent network locations in Active Directory, and placing domain controllers in their respective sites enables Active Directory to distinguish intra-site from intersite replication.

Next, create an additional site link containing Site1 and Site2. A site link defines the network path used by intersite replication and is where replication cost, availability, and schedule are configured. This dedicated link allows the Site1-to-Site2 replication schedule to be administered independently.

Finally, remove Site2 from DEFAULTIPSITELINK. The default site link can then retain the connection between Site1 and Site3, while the newly created link provides the separate Site1-to-Site2 path. As a result, each replication path has its own site-link schedule. This design also maintains direct connectivity from Site1 to both remote locations, minimizing replication interruptions and avoiding unnecessary reliance on DC2 as an intermediary for replication to DC3. Connection objects are not the appropriate control point for independently scheduling replication between sites; the schedule belongs to the site-link configuration. Microsoft documents site links and their use in controlling intersite replication at Designing the site topology and Understand default Active Directory site links.

QUESTION NO: 6

Your network contains a multi-site Active Directory Domain Services (AD DS) forest. Each Active Directory site is connected by using manually configured site links and automatically generated connections.

You need to minimize the convergence time for changes to Active Directory.

What should you do?

A. For each site link, modify the options attribute.

B. For each site link, modify the site link costs.

C. For each site link, modify the replication schedule.

D. Create a site link bridge that contains all the site links.

ANSWER: A

Explanation:

For each site link, modify the options attribute is correct because the site-link `options` attribute can enable intersite change notification. Specifically, setting the `USE_NOTIFY` flag causes a domain controller that receives a directory update to notify its intersite replication partner rather than waiting for the next normal intersite replication interval. The notified partner then requests the changes and can in turn notify downstream partners, substantially reducing the time required for an update to converge across the site topology.

This setting is appropriate when the WAN links between sites are sufficiently reliable and have enough available bandwidth to support more immediate replication traffic. It retains the existing site-link topology and the Knowledge Consistency Checker-generated connection objects, while changing the replication behavior associated with each site link. Administrators configure this through the site link's Options attribute in Active Directory Sites and Services or by using appropriate directory-management tooling. Microsoft describes intersite replication, site links, and the use of change notifications in its Active Directory replication concepts documentation. The underlying site-link configuration is represented by the Options attribute in the Active Directory schema.

QUESTION NO: 7

Your network contains an Active Directory Domain Services (AD DS) domain named contoso.com.

You have a domain-based Distributed File System (DFS) namespace named \\contoso.com\Public that is hosted on Server1.

You need to ensure that users can access the namespace if Server1 becomes unavailable.

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

NOTE: Each correct selection is worth one point.

A. Configure the namespace as a domain-based namespace.

B. Add Server2 as a namespace server.

C. Configure Server1 as a DFS Replication member.

D. Create a stand-alone namespace on Server2.

E. Enable Data Deduplication on the namespace servers.

ANSWER: A B

Explanation:

A domain-based namespace stores its configuration in Active Directory and can have multiple namespace servers. To provide namespace availability, add Server2 as an additional namespace server. Server2 then hosts a namespace target for the same domain-based namespace.

The namespace must be configured as a domain-based namespace because stand-alone namespaces do not use Active Directory for namespace configuration and do not support multiple namespace servers for fault tolerance. DFS Namespace availability applies to namespace referrals; availability of the shared folders themselves can require DFS Replication or another appropriate data-availability solution. For more information, see DFS Namespaces overview.

QUESTION NO: 8 - (DRAG DROP)

DRAG DROP

You have a server named Server1.

You plan to use Storage Spaces to expand the storage available to Server1. You attach eight physical disks to Server1. Four disks are HDDs and four are SSDs.

You need to create a volume on Server1 that will use the storage on all the new disks. The solution must provide the fastest read performance for frequently used files.

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

Select and Place:

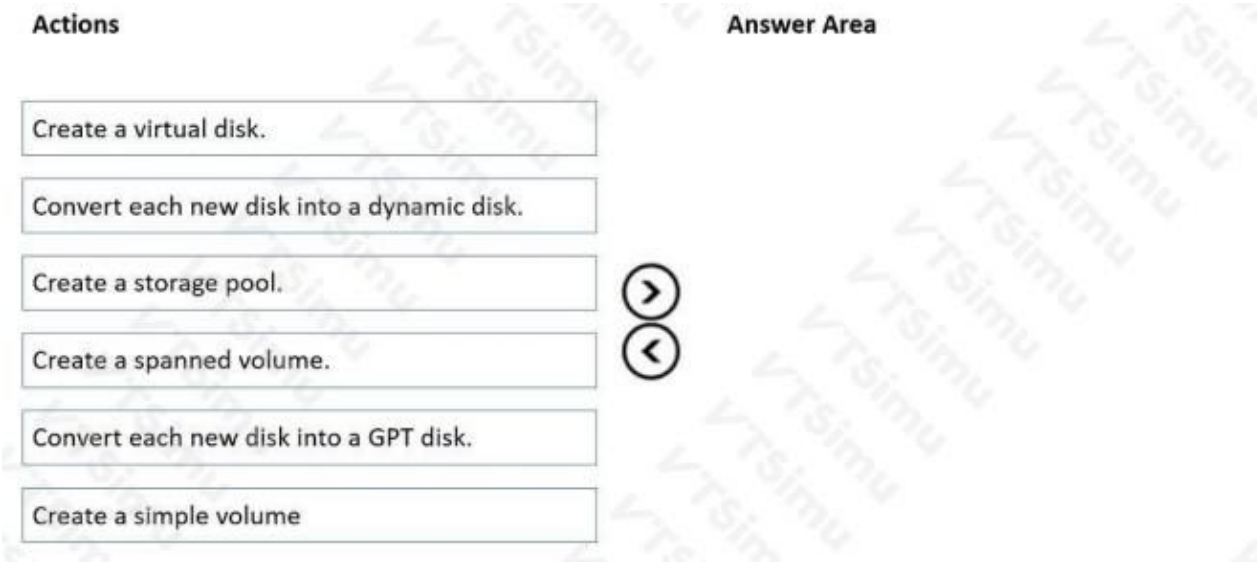

ANSWER:

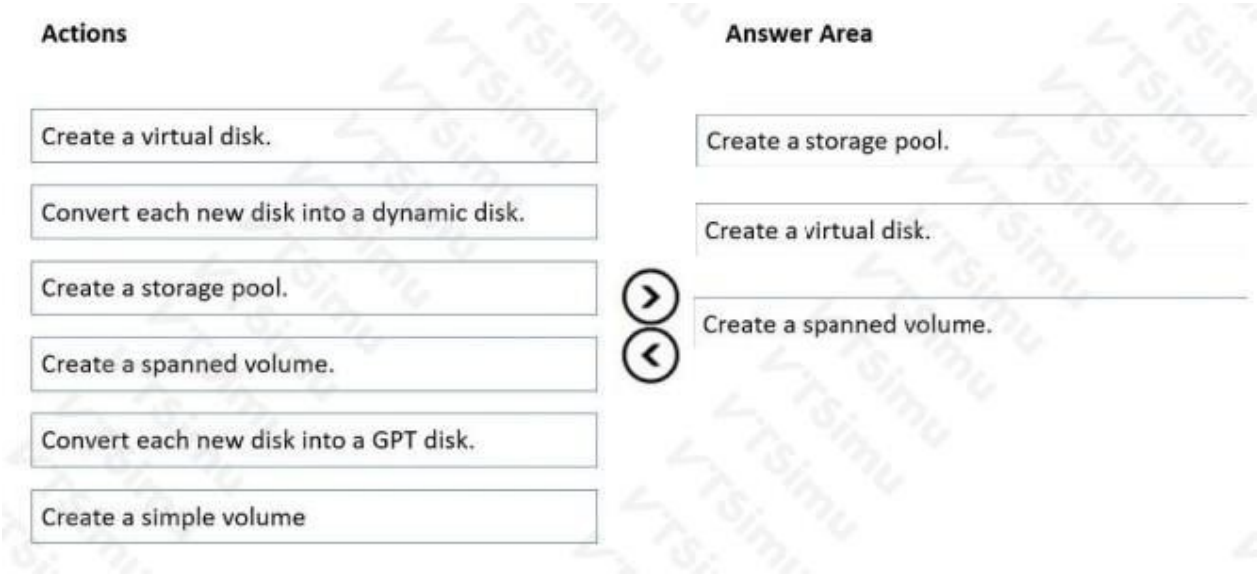

Explanation:

Storage Spaces begins by grouping the available physical drives into a storage pool. In this case, creating one pool from the four HDDs and four SSDs makes all of the newly attached storage available to Storage Spaces as a common set of resources. The pool is the required foundation for creating Storage Spaces virtual disks; individual disks do not need to be converted to dynamic disks for this process.

Next, create a virtual disk in the pool. A Storage Spaces virtual disk is the logical storage object built from capacity in the pool. Because the pool contains both SSD and HDD media, the virtual disk can be configured with storage tiers. Tiering places frequently accessed (“hot”) data on the SSD tier, where it benefits from substantially faster read performance, while less frequently used data is retained on the HDD tier. Storage Spaces automatically manages data movement between tiers based on usage patterns. This directly addresses the requirement to optimize read performance for frequently used files while using storage from all of the new disks.

After the virtual disk has been created, create a simple volume on it. The volume supplies the usable file system area: it can be assigned a drive letter or mount point and formatted, such as with NTFS or ReFS. This is the final step that exposes the Storage Spaces capacity for file storage on Server1. Microsoft’s Storage Spaces deployment workflow and tiering concepts are documented in Deploy standalone Storage Spaces and Storage Spaces storage tiers.

QUESTION NO: 9 - (HOTSPOT)

You have an onpremises DNS server named Server1 that runs Windows Server. Server 1 hosts a DNS zone named fabrikam.com

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

Name	Type	Description
Vnet1	Virtual network	Connects to the on-premises network by using a Site-to-Site VPN
VM1	Virtual machine	Runs Windows Server and has the DNS Server role installed
contoso.com	Private DNS zone	Linked to Vnet1
contoso.com	Public DNS zone	Contains the DNS records of all the platform as a service (PaaS) resources

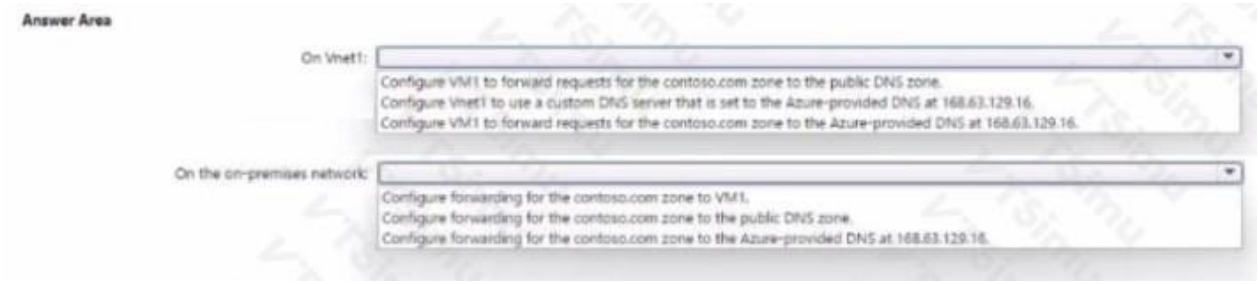

ANSWER:

Explanation:

To resolve records in an Azure Private DNS zone from an on-premises network, a DNS forwarder running in the Azure virtual network is required. In this design, VM1 provides that forwarder role. VM1 must be configured with a conditional forwarder for the private zone that sends queries to the Azure-provided DNS address, **168.63.129.16**. That address is the Azure recursive resolver endpoint used by resources inside an Azure virtual network to resolve names hosted in Azure Private DNS zones linked to that virtual network.

The on-premises DNS server cannot forward directly to 168.63.129.16, because this virtual IP is only available from within Azure virtual networks. Instead, Server1 forwards requests for the private DNS zone to VM1 by using VM1's reachable private IP address over the configured hybrid network connection. VM1 then forwards those requests internally to Azure-provided DNS, which can answer using the linked private DNS zone. This creates the required resolution path: on-premises clients query Server1, Server1 forwards the private-zone query to VM1, and VM1 forwards it to Azure DNS at 168.63.129.16.

Microsoft documents this DNS forwarder pattern for resolving Azure Private DNS zones from on-premises environments. See Azure Private DNS troubleshooting and resolution guidance and Azure Private DNS resolution for on-premises workloads using a DNS forwarder.

QUESTION NO: 10

You have a server named Server1 that hosts Windows containers.

You plan to deploy an application that will have multiple containers. Each container will be on the same subnet. Each container requires a separate MAC address and IP address. Each container must be able to communicate by using its IP address.

You need to create a Docker network that supports the deployment of the application.

Which type of network should you create?

A. NAT

B. transparent

C. l2bridge

D. l2tunnel

ANSWER: B

Explanation:

The **transparent** network driver is appropriate because it connects Windows containers directly to the physical network through an external Hyper-V virtual switch. Containers attached to this network receive connectivity on the same physical subnet as the host network and are addressable individually by their own IP addresses. This topology supports the requirement for each container to have a distinct MAC address and IP address that can be used for direct communication with other systems on that subnet.

With a transparent network, the network infrastructure, typically the subnet's DHCP service or statically assigned addressing, provides container IP configuration. The external virtual switch exposes the container endpoints to the network in a manner similar to separate networked devices. Therefore, systems on the same subnet can communicate with each container by using that container's assigned IP address, which is essential for applications whose individual container instances must be independently reachable.

Microsoft documents transparent networking as a Windows container network driver that provides direct connectivity to the physical network and supports assigning each container an address on the physical network. See Windows container network drivers and topologies and Windows container networking architecture.

QUESTION NO: 11

Your network contains a single-domain Active Directory Domain Services (AD DS) forest named conto.com. The forest contains the servers shown in the following exhibit table.

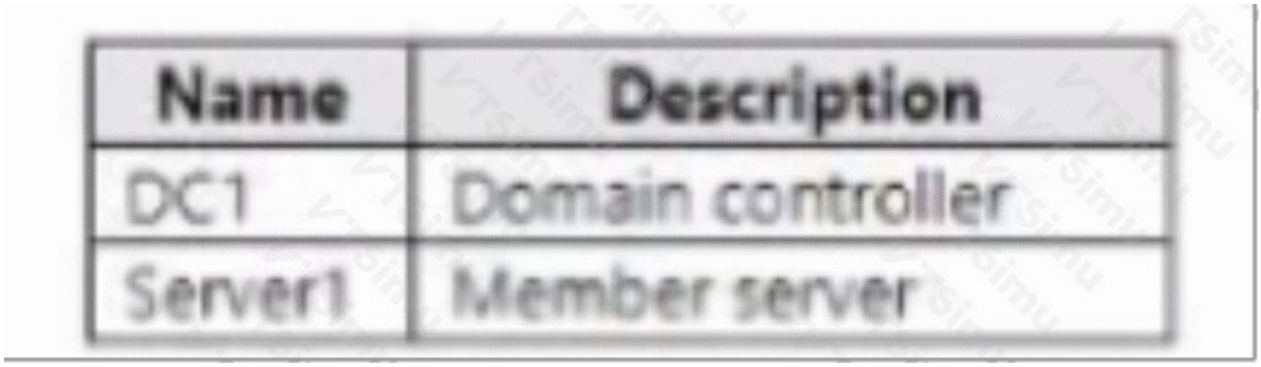

Name	Description
DC1	Domain controller
Server1	Member server

You plan to install a line-of-business (LOB) application on Server1. The application will install a custom windows services.

A new corporate security policy states that all custom Windows services must run under the context of a group managed service account (gMSA). You deploy a root key.

You need to create, configure, and install the gMSA that will be used by the new application.

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

NOTE: Each correct selection is worth one point.

A. On Server1, run the install-ADServiceAccount cmdlet.

B. On DC1, run the New-ADServiceAccount cmdlet.

C. On DC1, run the Set_ADComputer cmdlet.

D. ON DC1, run the Install-ADServiceAccount cmdlet.

E. On Server1, run the Get-ADServiceAccount cmdlet.

ANSWER: A B

Explanation:

Creating the gMSA requires running `New-ADServiceAccount`. This cmdlet creates the service-account object in Active Directory and can configure essential properties, including the DNS hostname, service principal names, and the computers or groups permitted to retrieve its managed password. The previously deployed Key Distribution Services root key is a

prerequisite that enables AD DS to generate and rotate the gMSA password automatically. In this scenario, the account should be configured so that Server1 is authorized to retrieve the password, either directly or through a group containing Server1.

After the account exists and Server1 has permission to retrieve its password, `Install-ADServiceAccount` must be run on Server1. Installation makes the managed service account available to the local computer and verifies that the computer can retrieve the account password from AD DS. The LOB application's Windows service can then be configured to use the gMSA logon identity, normally specified as `domain\account$`, without supplying or maintaining a password manually. Microsoft documents gMSA creation and its password-retrieval permissions in New-ADServiceAccount and documents local installation and validation in Install-ADServiceAccount.

QUESTION NO: 12

Your network contains an Active Directory Domain Services (AD DS) forest. The forest contains three domains. Each domain contains 10 domain controllers.

You plan to store a DNS zone in a custom Active Directory partition.

You need to create the Active Directory partition for the zone. The partition must replicate to only four of the domain controllers.

What should you use?

A. ntdsutil.exe

B. Active Directory Sites and Services

C. New-ADobject

D. Active Directory Administrative Center

ANSWER: A

Explanation:

ntdsutil.exe is the appropriate tool because it can create an Active Directory application directory partition, also called an application naming context. A custom DNS application directory partition is useful when DNS zone data must be replicated to a deliberately selected set of domain controllers rather than to every domain controller in a domain or throughout the forest. Using the `ntdsutil` Directory Service Management commands, an administrator creates the application partition and specifies the domain controller that hosts its first replica. Additional replica domain controllers can then be added explicitly, allowing the partition to be hosted and replicated on only the required four domain controllers. After the partition exists and the chosen DNS servers are enlisted for it, DNS can store an AD-integrated zone in that custom partition. This provides replication scope control while retaining Active Directory-integrated DNS storage and replication. Microsoft documents application directory partitions as directory partitions that can have replicas on a subset of domain controllers, and documents Ntdsutil as a command-line tool for managing AD DS, including application directory partitions. See Microsoft guidance on AD DS directory partitions and replication and Ntdsutil Directory Service Management reference.

QUESTION NO: 13

You have a DNS server named Server1 that runs Windows Server. Server1 hosts an Active Directory-integrated DNS zone named contoso.com.

You need to automatically remove stale resource records from contoso.com. The solution must allow records that are actively refreshed by clients to remain in the zone.

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

NOTE: Each correct selection is worth one point.

A. Enable aging for the contoso.com zone.

B. Enable automatic scavenging on Server1.

C. Configure DNSSEC signing for contoso.com.

D. Create a secondary zone for contoso.com.

E. Configure a conditional forwarder for contoso.com.

ANSWER: A B

Explanation:

DNS aging must be enabled on the zone so that dynamically registered records receive timestamps and become eligible for scavenging after the configured no-refresh and refresh intervals. Records refreshed by clients receive updated timestamps and are not removed as stale.

You must also enable automatic scavenging on the DNS server. The server-level scavenging process evaluates eligible records in zones configured for aging and removes records whose timestamps have exceeded the configured aging intervals. For more information, see Manage DNS zones and Aging and scavenging DNS records.

QUESTION NO: 14

Your network contains an Active Directory Domain Services (AD DS) forest. The forest contains three domains. Each domain contains 10 domain controllers.

You plan to store a DNS zone in a custom Active Directory partition.

You need to create the Active Directory partition for the zone. The partition must replicate to only four of the domain controllers.

What should you use?

A. DNS Manager

B. New-ADObject

C. dnscmd.exe

D. Windows Admin Center

ANSWER: C

Explanation:

dnscmd.exe is the appropriate tool because DNS application directory partitions are designed to provide a custom Active Directory-integrated DNS replication scope. Rather than using the built-in domain-wide or forest-wide DNS partitions, an administrator can create a named DNS directory partition and explicitly control which DNS servers host it. This permits the partition, and therefore a zone stored in it, to replicate only among the selected four domain controllers instead of all domain controllers in a domain or forest.

Use `dnscmd /CreateDirectoryPartition <partition-FQDN>` to create the custom partition. Then use `dnscmd /EnlistDirectoryPartition <partition-FQDN>` on each of the four selected DNS servers to enlist them as replicas. After the partition exists and the required servers are enlisted, create or change the AD-integrated zone so that it is stored in that custom directory partition. This approach provides the required granular replication control while retaining AD-integrated DNS storage and replication.

Microsoft documents the DNS command utility and its directory-partition operations at Dnscmd. Additional background on Active Directory application directory partitions is available in Understanding Active Directory site topology.

QUESTION NO: 15

You have an on-premises Active Directory Domain Services (AD DS) domain that syncs with an Azure Active Directory (Azure AD) tenant. You plan deploy 100 new Azure virtual machines that will run Windows Server. You need to ensure that each new virtual machine is joined to the AD DS domain. What should you use?

A. Azure AD Connect

B. a Group Policy Object (GPO)

C. an Azure Resource Manager (ARM) template

D. an Azure management group

ANSWER: C

Explanation:

an Azure Resource Manager (ARM) template is correct because it provides a consistent, repeatable way to deploy all 100 Windows Server virtual machines and configure domain joining as part of each deployment. The template can deploy the virtual machine and include the Azure VM extension for joining an existing AD DS domain, commonly the JsonADDomainExtension. Extension settings specify the domain name, organizational unit path if required, and a domain account with permission to join computers; sensitive credentials should be supplied securely through protected settings or Azure Key Vault rather than stored in plaintext. For this to work, the Azure virtual machines must also have network connectivity to domain controllers and use DNS servers that can resolve the AD DS domain. An ARM template is therefore well suited to a large-scale, standardized provisioning task, ensuring each newly deployed VM receives the same domain-join configuration automatically. Microsoft documents domain joining through the AD domain extension and supports defining VM extensions in Azure Resource Manager deployments. See Microsoft's AD domain join extension documentation and Microsoft's ARM template overview.

QUESTION NO: 16 - (HOTSPOT)

You have on-premises file servers that run Windows Server as shown in the following table.

Name	Relevant folder
Server1	D:\Folder1, E:\Folder2
Server2	D:\Data

You have the Azure file shares shown in the following table.

Name	Location
share1	East US
share2	East US

You add a Storage Sync Service named Sync1 and an Azure File Sync sync group named Group1. Group1 uses share1 as a cloud endpoint.

You register Server1 and Server2 with Sync1. You add D:\Folder1 from Server1 as a server endpoint in Group1.

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
You can add share2 as a cloud endpoint in Group1.	○	○
You can add E:\Folder2 from Server1 as a server endpoint in Group1.	○	○
You can add D:\Data from Server2 as a server endpoint in Group1.	○	○

ANSWER:

Explanation:

An Azure File Sync sync group is the synchronization boundary that connects server endpoints to an Azure file share. Its cloud side is deliberately limited to one cloud endpoint: the Azure file share selected when the sync group is configured. In this scenario, Group1 already uses share1 as that cloud endpoint. The sync group therefore has its single cloud endpoint established, while additional Azure file shares require their own separate sync groups. Microsoft documents this sync-group design in its Azure File Sync deployment guide.

Server endpoints are more flexible. A sync group can include multiple server endpoints, including endpoints hosted on the same registered server or on different registered servers. The important path requirement is that endpoints on a server must not overlap within the same sync group. D:\Folder1 and E:\Folder2 are separate folders on separate volumes on Server1, so they do not overlap and can both participate in Group1. Adding D:\Data from Server2 is also supported because Server2 is registered with Sync1 and Azure File Sync is designed to synchronize data between multiple server endpoints through the sync group's cloud endpoint.

Accordingly, the correct selections are No for adding share2 as another cloud endpoint in Group1, Yes for adding E:\Folder2 from Server1 as a server endpoint, and Yes for adding D:\Data from Server2 as a server endpoint. The broader endpoint model is described in Microsoft's Azure File Sync overview, which explains that a sync group contains one cloud endpoint and one or more server endpoints.

QUESTION NO: 17

You have an on premises Active Directory Domain Services (AD DS) domain that syncs with an Azure Active Directory (Azure AD) tenant.

Yo u plan to implement self-service password reset (SSPR) in Azure AD.

You need to ensure that users that reset their passwords by using SSPR can use the new password resources in the AD DS domain.

What should you do?

A. Deploy the Azure AD Password Protec tion proxy service to the on premises network.

B. Run the Microsoft Azure Active Directory Connect wizard and select Password writeback.

C. Grant the Change password permission for the domain to the Azure AD Connect service account.

D. Grant the impersonat e a client after authentication user right to the Azure AD Connect service account.

ANSWER: B

Explanation:

Run the Microsoft Azure Active Directory Connect wizard and select Password writeback. Password writeback is the Microsoft Entra Connect feature that securely sends a password changed or reset in the cloud back to the on-premises Active Directory Domain Services environment. This is required for hybrid users who use self-service password reset and must subsequently authenticate to resources that rely on their AD DS credentials, such as domain-joined servers, file shares, or other on-premises applications.

After password writeback is enabled in the Microsoft Entra Connect wizard, the service uses the configured synchronization connection and the required directory permissions to apply eligible cloud password resets to the corresponding on-premises user account. The user's AD DS password is therefore updated to match the new password established through SSPR. Password writeback is available with Microsoft Entra ID licensing that supports the feature, and it must be enabled both in the tenant's SSPR configuration and in Microsoft Entra Connect. Microsoft documents the password writeback prerequisites and configuration process at Enable self-service password reset writeback and describes the hybrid password-reset capability at Password writeback.

QUESTION NO: 18

You have a server named Server1 that runs Windows Server 2019 and hosts a container named Contained. Contained uses a Windows Server 2019 base image that was built by using a Docker file.

You upgrade Server1 to Windows Serve r 2022.

You need to ensure that Contained will run on Server1. The solution must minimize administrative effort.

What should you do?

A. Start Contained in process isolation mode.

B. Modify the Docker file.

C. Start Contained in Hyper-V isolation mode.

D. R ebuild the base image for Contained.

ANSWER: C

Explanation:

Start Contained in Hyper-V isolation mode. Windows containers using process isolation share the host kernel, so their base-image version must meet the applicable host and container version-compatibility requirements. After Server1 is upgraded to Windows Server 2022, the existing Windows Server 2019 container image does not provide the desired process-isolation compatibility path. Hyper-V isolation addresses this by running the container inside a lightweight utility virtual machine that supplies a kernel matching the container image. Therefore, the Windows Server 2019-based container can continue to run on the Windows Server 2022 host without rebuilding its image or changing the Dockerfile. This is the lowest-effort solution because the existing image and container definition can be retained; the isolation mode is selected when the container is started, for example by using Docker's `--isolation=hyperv` setting. Microsoft documents that Hyper-V isolation provides broader version compatibility by pairing the container with a matching kernel, while preserving the container deployment model. See Windows container version compatibility and Hyper-V isolation for Windows containers.

QUESTION NO: 19 - (SIMULATION)

Task 8

You plan to delegate the management of a ONS zone named fabnkam.com located on DO to the BranchAdmins grou p. You need to ensure that you can grant permissions to the fabikam.com zone.

ANSWER: See the explanation for the answer

Explanation:

On **DC1**, make the **fabikam.com** zone Active Directory-integrated (if it is not already), because DNS zone ACLs can be assigned only to AD-integrated zones.

1. Open **DNS Manager** (`dnsmgmt.msc`).
2. Expand the server, select **Forward Lookup Zones**, right-click **fabikam.com**, and select **Properties**.
3. On the **General** tab, click **Change** beside *Type*, select **Store the zone in Active Directory (available only if DNS server is a writable domain controller)**, and select **Primary zone**. Confirm the appropriate AD replication scope.
4. Open the zone **Properties** again and select the **Security** tab.
5. Click **Add**, add `BranchAdmins`, and resolve the group name.
6. Select `BranchAdmins` and grant **Allow: Full Control** (or, at minimum, Read, Write, Create All Child Objects, and Delete All Child Objects required for full record management).
7. Click **Apply** and **OK**.

This delegates management of DNS records in the **fabikam.com** zone to the `BranchAdmins` group without making its members DNS server administrators.

QUESTION NO: 20

You need to implement a name resolution solution that meets the networking requirements. 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 private DNS zone named corp.fabhkam.com.

B. Create a virtual network link in the coip.fabnkam.c om Azure private DNS zone.

C. Create an Azure DNS zone named corp.fabrikam.com.

D. Configure the DNS Servers settings for Vnet1.

E. Enable autoregistration in the corp.fabnkam.com Azure private DNS zone.

F. On DC3, install the DNS Server role.

G. Configure a conditional forwarder on DC3.

ANSWER: D F

Explanation:

Configure the DNS Servers settings for Vnet1 and install the DNS Server role on DC3. In a hybrid Active Directory Domain Services deployment, Azure virtual machines that must locate domain controllers and resolve records in an AD-integrated DNS zone need to use DNS servers that host, or can resolve, the AD DS DNS namespace. Installing the DNS Server role on DC3 enables that domain controller to provide DNS service and host the AD-integrated DNS data required for the domain. Domain controllers commonly run DNS because AD DS relies on DNS service-location records to identify domain controllers, Kerberos services, Global Catalog servers, and other directory services.

Configuring the DNS Servers settings for Vnet1 directs virtual machines in that virtual network to use DC3's private IP address as their custom DNS server. This is essential because Azure-provided DNS does not host AD-integrated zones or dynamically register AD DS records. After changing virtual-network DNS settings, affected virtual machines may need to renew their DHCP lease or be restarted so that they receive the updated DNS server configuration. Microsoft documents DNS requirements for AD DS in Azure and the use of custom DNS servers for virtual networks at Extend Active Directory Domain Services to Azure and Name resolution for Azure resources.

QUESTION NO: 21

Your network contains an Active Directory Domain Services (AD DS) forest named contoso.com.

You plan to allow administrators to restore deleted user accounts and organizational units without restoring domain controllers from backup.

You need to enable the required Active Directory feature for the forest.

What should you do?

A. Enable the Active Directory Recycle Bin optional feature.

B. Enable the AD DS auditing policy.

C. Create a custom Active Directory application partition.

D. Configure a global catalog server in each site.

ANSWER: A

Explanation:

Enable the **Active Directory Recycle Bin** optional feature. When enabled, the Active Directory Recycle Bin preserves deleted objects and most of their attributes, which allows administrators to restore deleted users, groups, and organizational units without performing an authoritative restore from backup.

The feature is enabled at the forest scope by using Active Directory Administrative Center or the `Enable-ADOptionalFeature` PowerShell cmdlet. Enabling the Recycle Bin is irreversible, so the forest functional level and operational requirements should be verified before enabling it. For more information, see Active Directory Recycle Bin.

QUESTION NO: 22

You have a two-node failover cluster named Cluster1. The cluster nodes are located in an on-premises datacenter.

You have an Azure subscription. Cluster1 does not have shared storage that can be used as a disk witness.

You need to configure a witness that helps Cluster1 maintain quorum if one node becomes unavailable. The solution must minimize on-premises infrastructure requirements.

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 Storage account.

B. Configure Cloud Witness in the quorum settings for Cluster1.

C. Add a disk witness to Cluster1.

D. Create a Cluster Shared Volume.

E. Configure a file share witness on one of the cluster nodes.

ANSWER: A B

Explanation:

A cloud witness uses an Azure Storage account as the quorum witness location. You must create or identify an Azure Storage account that the cluster can use, and then configure the cluster quorum settings to use Cloud Witness with the storage account information.

Cloud Witness is appropriate when a cluster has internet connectivity and an Azure subscription because it avoids deploying and maintaining a third on-premises server or file share solely for quorum. The cluster communicates with Azure Storage by using HTTPS. For more information, see Deploy a cloud witness for a failover cluster.

QUESTION NO: 23

You have two servers named Host1 and Host2 that run Windows Server and have the Hyper-V server role installed. Both servers are members of an Active Directory Domain Services (AD DS) domain.

You plan to perform live migrations between Host1 and Host2 by using Hyper-V Manager from a management computer named Server1.

You need to ensure that administrators can initiate the live migrations from Server1 without providing credentials for the destination host.

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

NOTE: Each correct selection is worth one point.

A. Configure Kerberos as the live migration authentication protocol on Host1 and Host2.

B. Configure constrained delegation for Host1 and Host2 to the Microsoft Virtual System Migration Service.

C. Configure CredSSP as the live migration authentication protocol on Host1 and Host2.

D. Add the Server1 computer account to the Hyper-V Administrators group on Host1 and Host2.

E. Enable enhanced session mode on Host1 and Host2.

ANSWER: A B

Explanation:

You should configure Kerberos authentication for live migration and configure constrained delegation for the Hyper-V hosts. Kerberos supports credential delegation for a migration initiated remotely from Server1. CredSSP does not support this double-hop scenario because it requires that the migration be initiated from the source host.

In Active Directory Users and Computers, configure the computer accounts for Host1 and Host2 to trust each other for delegation to the Microsoft Virtual System Migration Service. This permits each host to delegate an administrator's Kerberos credentials to the other host only for the live-migration service. For more information, see Set up hosts for live migration without Failover Clustering.

QUESTION NO: 24

You need to meet the technical requirements for VM3.

On which volumes can you enable Data Deduplication?

A. C and D only

B. D only

C. C, D, E, and F

D. D and E only

E. D, E, and F only

ANSWER: D

Explanation:

D and E only is correct. Data Deduplication is enabled on individual data volumes rather than globally for an entire server. To be eligible, a volume must meet the feature's supported file-system and volume requirements and must not be the operating-system volume. In the VM3 volume configuration, volumes D and E are the volumes that meet those requirements, so they can be selected as deduplication targets. After the Data Deduplication role service is installed, an administrator can enable the feature by using Server Manager or the `Enable-DedupVolume` PowerShell cmdlet and then select an

appropriate usage type, such as General purpose, Hyper-V, or Backup. Deduplication processes data after it is written and maintains the volume's logical file layout while storing duplicate data only once, thereby reducing storage consumption. Microsoft also documents workload-specific guidance and interoperability considerations, particularly when deduplicating volumes used for virtual-machine storage. See the Microsoft guidance for installing and enabling Data Deduplication and the Data Deduplication interoperability documentation.

QUESTION NO: 25

You have an on-premises Active Directory Domain Services (AD DS) domain that syncs with an Azure Active Directory (Azure AD) tenant The on-premises network is connected to Azure by using a Site-to-Site VPN. You have the DNS zones shown in the following table.

Name	Location	Description
contoso.com	A domain controller named DC1 on the on-premises network	Provides name resolution on-premises
fabrikam.com	An Azure private DNS zone	Provides name resolution for all Azure virtual networks

You need to ensure that names from (aDiifcam.com can be resolved from the on-premises network 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 conditional forwarder for fabrikam.com on DC1.

B. Create a stub zone for fabrikam.com on DC1.

C. Create a secondary zone for fabnlcam.com on DO.

D. Deploy an Azure virtual machine that runs Windows Server. Modify the DNS Servers settings for the virtual network.

E. Deploy an Azure virtual machine that runs Windows Server. Configure the virtual machine &s a DNS forwarder.

ANSWER: A E

Explanation:

Creating a conditional forwarder for fabrikam.com on DC1 directs DNS queries for that specific namespace from the on-premises DNS environment to a designated DNS server, rather than attempting resolution through the normal recursive path. This provides the required targeted resolution path for clients in the on-premises AD DS domain.

Deploying an Azure virtual machine running Windows Server and configuring it as a DNS forwarder provides the DNS endpoint reachable through the site-to-site VPN. The forwarder can accept the forwarded fabrikam.com queries from DC1 and forward them to Azure's internal DNS service, which can resolve records hosted in an Azure Private DNS zone. This is the supported pattern when on-premises workloads must resolve Azure private DNS names: an on-premises conditional forwarder sends requests to a DNS forwarder in Azure, and that forwarder uses Azure-provided DNS for private-zone resolution.

Configure the DNS forwarder with connectivity and firewall rules that allow DNS traffic from the on-premises DNS server across the VPN. For implementation guidance, see Azure Private Endpoint DNS configuration for on-premises workloads and Azure hybrid DNS resolution.

QUESTION NO: 26

Your network contains an Active Directory Domain Services (AD DS) domain named contoso.com.

A DNS server named DNS1 hosts the contoso.com zone. A separate DNS server named DNS2 hosts a zone named branch.contoso.com.

You need to ensure that clients that use DNS1 can resolve names in the branch.contoso.com namespace.

What should you create on DNS1?

A. a delegation for branch.contoso.com

B. a reverse lookup zone for branch.contoso.com

C. a stub zone for contoso.com

D. a root hint for DNS2

ANSWER: A

Explanation:

You should create a **delegation** for branch.contoso.com in the contoso.com zone. A delegation identifies the authoritative DNS servers for a child zone by creating NS records, and normally associated glue A or AAAA records, in the parent zone. DNS1 can then refer clients to DNS2 for records in branch.contoso.com.

A conditional forwarder is useful when a DNS server must forward queries for another namespace, but a delegation is the appropriate DNS hierarchy mechanism when DNS2 hosts an authoritative child zone of contoso.com. For more information, see Delegate a DNS zone.

QUESTION NO: 27

You have an Azure virtual machine named VM1 that runs Windows Server.

You perform the following actions on VM1:

- Create a folder named Folder1 on volume C.
- Create a folder named Folder2 on volume D.
- Add a new data disk to VM1 and create a new volume that is assigned drive letter

E. - Install an app named App1 on volume E.

You plan to resize VM1.

Which objects will present after you resize VM1?

A. Folder1, volume E, and App1 only

B. Folder1 only

C. Folder1 and Folder2 only

D. Folder1, Folder2, App1, and volume E

E. Install an app named App1 on volume E.
You plan to resize VM1.
Which objects will present after you resize VM1?

ANSWER: A

Explanation:

Folder1, volume E, and App1 only is correct. The C: drive is the virtual machine's operating-system disk, which is persistent Azure managed storage. Therefore, Folder1 remains available after the virtual machine is resized. The disk added to create E: is a persistent data disk. Resizing a VM changes its compute size but does not remove attached managed data disks, so the volume on that disk and the application installed on it remain present after the resize operation.

In contrast to persistent OS and data disks, the local temporary disk supplied with many Azure VM sizes is not intended for durable data. Azure can reinitialize that disk during maintenance events, redeployment, resizing, or other host changes. Its contents must therefore not be relied on for application data or files that need to survive infrastructure operations. Microsoft recommends storing durable data on managed disks or other persistent Azure storage services instead.

For details, see Azure managed disks overview and Temporary disk guidance for Azure virtual machines.

QUESTION NO: 28 - (SIMULATION)

Task 9

You need to ensure that all the computers in the domai n use DNSSEC to resolve names in the adatum.com zone.

ANSWER: See the explanation for the answer

Explanation:

Configure a Name Resolution Policy Table (NRPT) rule by using a domain-linked Group Policy Object. This forces domain DNS clients to request DNSSEC validation for the ADatum namespace; signing the zone alone does not make every client require DNSSEC.

1. Open **Group Policy Management** and create or edit a GPO linked to the **adatum.com** domain (so it applies to all domain computer accounts).
2. Navigate to:
   `Computer Configuration > Policies > Windows Settings > Name Resolution Policy`
3. Right-click **Name Resolution Policy** and select **New**.
4. Set the namespace to `.adatum.com` (or `adatum.com`, depending on the UI's accepted namespace format).
5. On the DNSSEC settings for the rule, select **Enable DNSSEC in this rule**.
6. Select **Require DNS clients to check that name and address data has been validated**.
7. Save the rule and allow Group Policy to refresh (or run `gpupdate /force` on clients).

The DNS servers used by the clients must already be DNSSEC-capable and able to validate the signed `adatum.com` zone. The required client-wide configuration is the NRPT rule with DNSSEC enabled and validation required.

QUESTION NO: 29

You need to implement an availability solution for DHCP that meets the networking requirements.

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

NOTE: Each correct selection is worth one point.

A. On DHCP1. create a scope that contains 25 percent of the IP addresses from Scope2.

B. On the router in each office, configure a DHCP relay.

C. DHCP2. configure a scope that contains 25 percent of the IP addresses from Scope 1 .

D. On each DHCP server, install the Failover Clustering feature and add the DHCP cluster role.

E. On each DHCP scope, configure DHCP failover.

ANSWER: B E

Explanation:

Configuring DHCP failover on each DHCP scope provides the required DHCP service availability. A DHCP failover relationship uses two DHCP servers as partners and synchronizes scope configuration and lease information between them. It can operate in load-balance mode, where both servers service clients, or hot-standby mode, where a standby partner takes responsibility when the active server is unavailable. This preserves lease data and enables clients to continue obtaining or renewing addresses during a server outage. The relationship is configured per scope, so every scope that requires availability must be included in DHCP failover.

Configuring a DHCP relay on the router in each office ensures that clients can reach the DHCP servers when the clients and servers are on different IP subnets. DHCP client discovery begins as a local broadcast, which routers do not normally

forward. A relay agent receives that broadcast and forwards it as a unicast request to the configured DHCP server addresses. For a resilient design, the relay configuration should forward requests to both failover partners, allowing either available server to respond. Together, DHCP failover protects the service and lease database, while DHCP relay enables cross-subnet client access. Microsoft documents DHCP failover at DHCP failover and relay-agent deployment at Deploy a DHCP relay agent.

QUESTION NO: 30

You need to implement a name resolution solution that meets the networking requirements.

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

NOTE: Each correct selection is worth one point.

A. Configure the DNS Servers settings for Vnet1.

B. On DC3, install the DNS Server role.

C. Create a virtual network link in the corp.fabrikam.com Azure private DNS zone.

D. Configure a conditional forwarder on DC3.

E. Enable autoregistration in the corp.fabrikam.com Azure private DNS zone.

F. Create an Azure private DNZ zone named corp.fabrikam.com.

G. Create an Azure DNZ zone named corp.fabrikam.com.

ANSWER: A B

Explanation:

Installing the DNS Server role on DC3 makes DC3 able to host DNS zones and answer DNS queries for the required namespace. The DNS Server role is the Windows Server component that provides DNS name-resolution services, so it must be present before DC3 can function as the designated DNS server for client virtual machines. Microsoft documents deployment of this role through Server Manager or Windows PowerShell in its DNS Server installation and configuration guidance.

Configuring the DNS Servers settings for Vnet1 then directs Azure virtual machines connected to that virtual network to use the IP address of DC3 for DNS resolution. Azure virtual machines obtain DNS configuration from the virtual network's DNS settings; selecting custom DNS servers is therefore the required mechanism for centrally assigning DC3 as their resolver. After changing virtual-network DNS settings, virtual machines may need to be restarted or have their DHCP lease renewed before the updated DNS configuration is applied. Microsoft describes this custom-DNS behavior in Name resolution for resources in Azure virtual networks.

QUESTION NO: 31

What sho uld you implement for the deployment of DC3?

A. Azure Active Directory Domain Services (Azure AD DS}

B. Azure AD Application Proxy

C. an Azure virtual machine

D. an Azure AD administrative unit

ANSWER: C

Explanation:

an Azure virtual machine is the appropriate implementation for deploying DC3 when DC3 is intended to be an additional Active Directory Domain Services domain controller. A domain controller runs the AD DS role on Windows Server and must be joined to the existing domain, promoted as a domain controller, and configured to replicate directory data with the other domain controllers. An Azure virtual machine provides the required Windows Server operating system, administrative control, network connectivity to the existing AD DS environment, and persistent storage needed for this deployment. In a hybrid design, placing an additional domain controller in an Azure virtual network can also provide directory-service availability for Azure-hosted workloads, provided that virtual network DNS settings, connectivity, and site configuration support AD DS replication. Microsoft documents deploying AD DS domain controllers on Azure virtual machines and recommends using multiple domain controllers for resiliency. See Extend Active Directory Domain Services to Azure and Deploying AD DS domain controllers.

QUESTION NO: 32

You need to implement an availability solution for DHCP that meets the networking requirements.

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

NOTE: Each correct selection is worth one point.

A. On DHCP1, create a scope that contains 25 percent of the IP addresses from Scope2.

B. On the router in each office, configure a DHCP relay.

C. DHCP2, configure a scope that contains 25 percent of the IP addresses from Scope1.

D. On each DHCP server, install the Failover Clustering feature and add the DHCP cluster role.

E. On each DHCP scope, configure DHCP failover.

ANSWER: B E

Explanation:

Configuring DHCP failover on each DHCP scope provides the DHCP high-availability mechanism. DHCP failover establishes a relationship between two DHCP servers that replicate lease information and scope configuration for the protected scopes. Consequently, either partner can continue responding to clients if the other DHCP server becomes unavailable. The relationship can be configured in load-balance mode for active use of both servers, or hot standby mode where a standby server assumes service when required. It is configured per scope, so all scopes that require availability must be included in the failover relationship.

Configuring a DHCP relay on the router in each office enables client DHCP requests to reach the DHCP servers when the clients and servers reside on different IP subnets. DHCP client discovery traffic is broadcast traffic and is not routed between subnets by default. A relay agent receives the client broadcast, forwards it as unicast traffic to the configured DHCP servers, and supplies subnet information so the servers can allocate an address from the appropriate scope. Together, DHCP failover protects DHCP service continuity, while DHCP relay makes the centralized, redundant DHCP service reachable from every office subnet. See DHCP failover in Windows Server and Deploy a DHCP relay agent.

QUESTION NO: 33

You have an on-premises server named Server1 that runs Windows Server. Server1 is connected to Azure Arc.

You need to collect events from the System event log on Server1 in a Log Analytics workspace. The solution must use the Azure Monitor agent.

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

NOTE: Each correct selection is worth one point.

A. Deploy the Azure Monitor agent extension to Server1.

B. Create a data collection rule that collects System event log events and associate the rule with Server1.

C. Install the Log Analytics agent on Server1.

D. Enable Microsoft Sentinel on the Log Analytics workspace.

E. Create an Azure Automation runbook.

ANSWER: A B

Explanation:

You must deploy the **Azure Monitor agent** extension to Server1. For Azure Arc-enabled servers, Azure extensions provide the mechanism to install and manage the Azure Monitor agent on the non-Azure machine.

You must also create and associate a **data collection rule** with Server1. The data collection rule specifies the Windows event log source, such as the System log, and the destination Log Analytics workspace. The agent uses the associated rule to determine which data to collect and where to send it. For more information, see Manage the Azure Monitor Agent and Data collection rules in Azure Monitor.