Microsoft.AZ-800.v2026-02-17.q122

Exam Code:	AZ-800
Exam Name:	Administering Windows Server Hybrid Core Infrastructure
Certification Provider:	Microsoft
Free Question Number:	122
Version:	v2026-02-17
# of views:	112
# of Questions views:	1220
https://www.freepdfdumps.com/Microsoft.AZ-800.v2026-02-17.q122.html	

NEW QUESTION: 1

You have a server named Server1 that runs Windows Server and contains three volumes named C, D, and E.

Files are stored on Server1 as shown in the following table.

Name	In volume	Size
File1	C	500 KB
File2	D	10 KB
File3	D	1 MB

For volume D, Data Deduplication is enabled and set to General purpose file server.

You perform the following actions:

* Move File1 to volume D.
* Copy File2 to volume D and name the copy File4.
* Move File3 to volume E

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

NOTE: Each correct selection is worth one point.

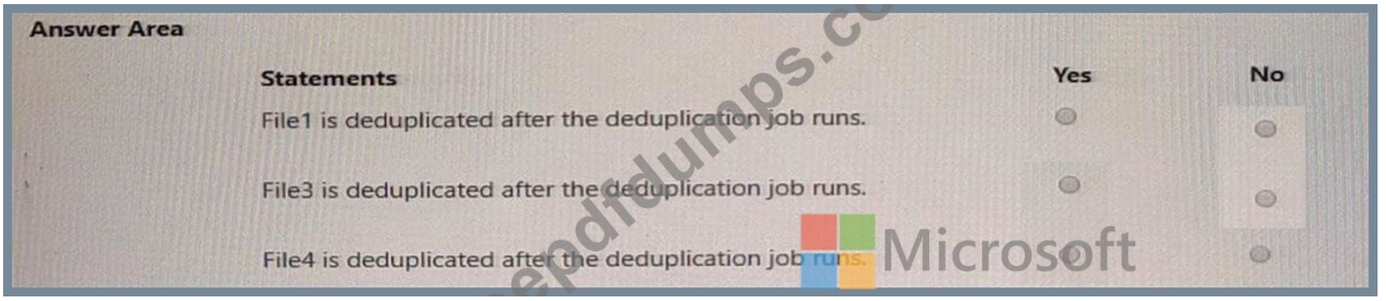

Answer:

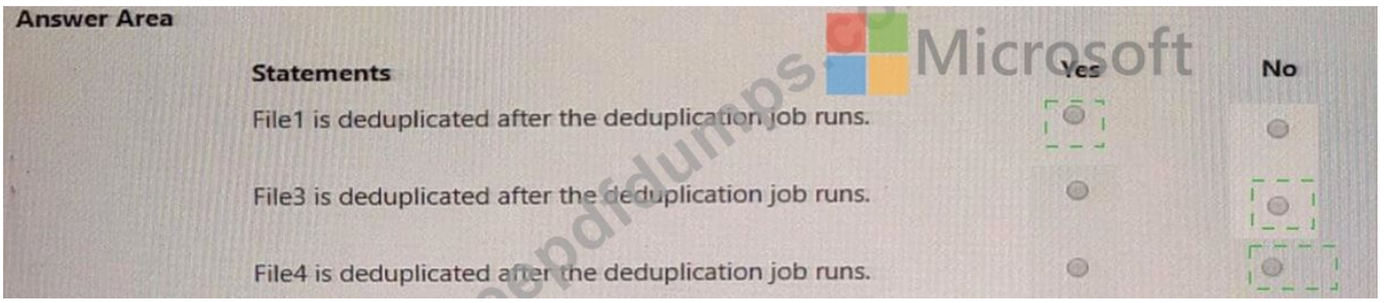

Explanation:

< File1 is deduplicated after the deduplication job runs. - YES

File3 is deduplicated after the deduplication job runs. - NO

File4 is deduplicated after the deduplication job runs. - NO

The Administering Windows Server Hybrid Core Infrastructure materials explain that Data Deduplication operates per-volume and only processes files on volumes where the role is enabled. The guide states that deduplication "is applied only to NTFS/ReFS volumes on which the Data Deduplication role is enabled," and that the General-purpose file server usage type applies default policies for typical data shares. It further specifies the file size limits: "Files smaller than 32 KB are not deduplicated; supported files are 32 KB up to multiple terabytes," and clarifies that optimization jobs process eligible files during scheduled runs.

Applying those rules:

* Volume D has Data Deduplication enabled (General-purpose). After moving File1 (500 KB) from C: to D:, it resides on a deduplicated volume and exceeds the minimum size threshold, so it will be deduplicated by the next optimization job.
* File3 (1 MB) is moved off the deduplicated volume (to E:), and dedup only affects enabled volumes; therefore it will not be deduplicated.
* File4 is a copy of File2 (10 KB) on D:. Because the file is smaller than the 32-KB minimum, it is not deduplicated even though it is on a deduplicated volume.

Thus the correct outcomes are YES for File1, NO for File3, and NO for File4.

NEW QUESTION: 2

You create a new Azure subscription.

You plan to deploy Azure Active Directory Domain Services (Azure AD DS) and Azure virtual machines.

The virtual machines will be joined to Azure AD DS.

You need to deploy Active Directory Domain Services (AD DS) to ensure that the virtual machines can be deployed and joined to Azure AD DS.

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.

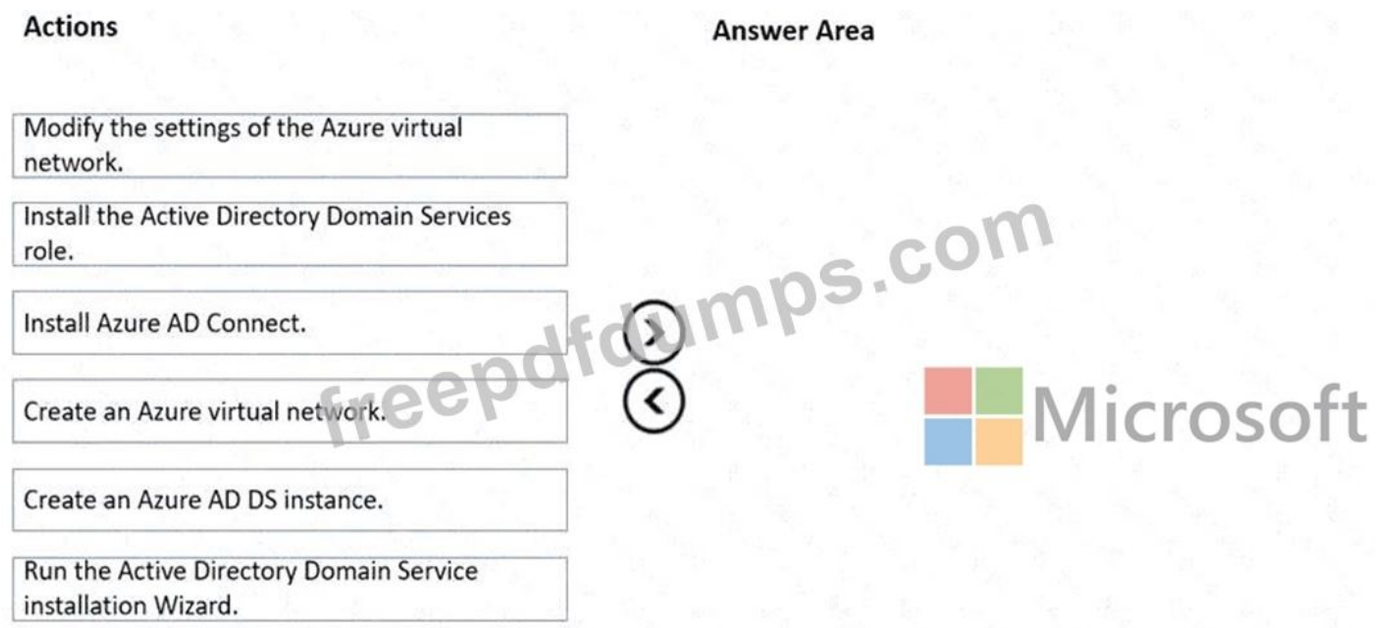

Answer:

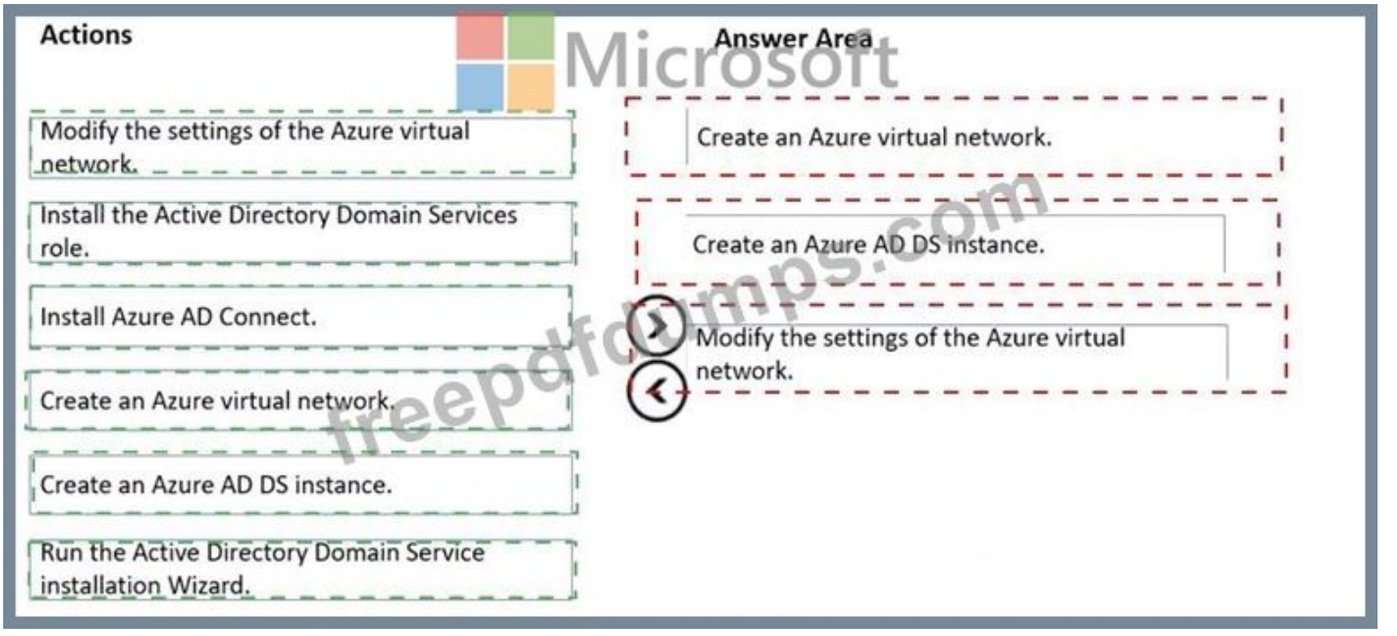

Explanation:

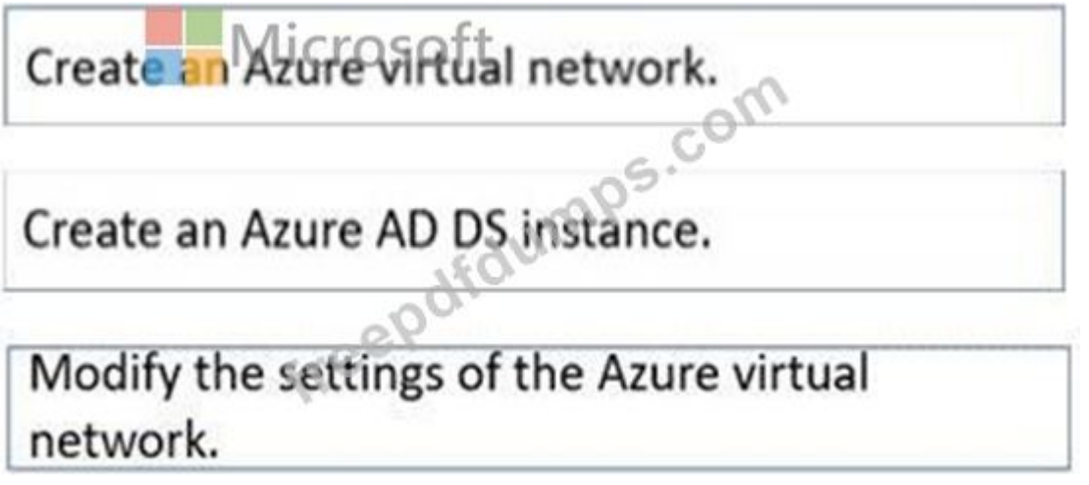

Reference:

The Administering Windows Server Hybrid Core Infrastructure guidance for deploying Azure Active Directory Domain Services (Azure AD DS) in a new subscription is explicit about the required sequence.
First, you prepare networking: "Before you enable Azure AD DS, create or select an Azure virtual network and subnet that will host the managed domain." Next, you provision the managed domain: "Enable Azure AD DS into the chosen virtual network/subnet; the service deploys managed domain controllers and exposes domain IP addresses." Finally, you update DNS for the VNet so VMs can locate the domain: "After the managed domain is provisioned, configure the virtual network DNS servers to the IP addresses of the Azure AD DS domain controllers so that virtual machines can resolve the domain and join it." The course also clarifies what you do not do with Azure AD DS: "Azure AD DS is a managed domain; you do not install the AD DS role or run the AD DS installation wizard on your own VMs." And: "Azure AD Connect is only required when synchronizing identities from an on-premises AD; it isn't required for a cloud- only deployment of Azure AD DS." Therefore, the correct sequence to ensure VMs can be deployed and joined to Azure AD DS is:
Create an Azure virtual network,
Create an Azure AD DS instance (into that VNet/subnet),
Modify the VNet DNS settings to point to the managed domain IPs.

NEW QUESTION: 3
You have an Azure Active Directory Domain Services (Azure AD DS) domain.
You create a new user named Admin1.
You need Admin1 to deploy custom Group Policy settings to all the computers in the domain. The solution must use the principle of least privilege.
What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point

Answer:

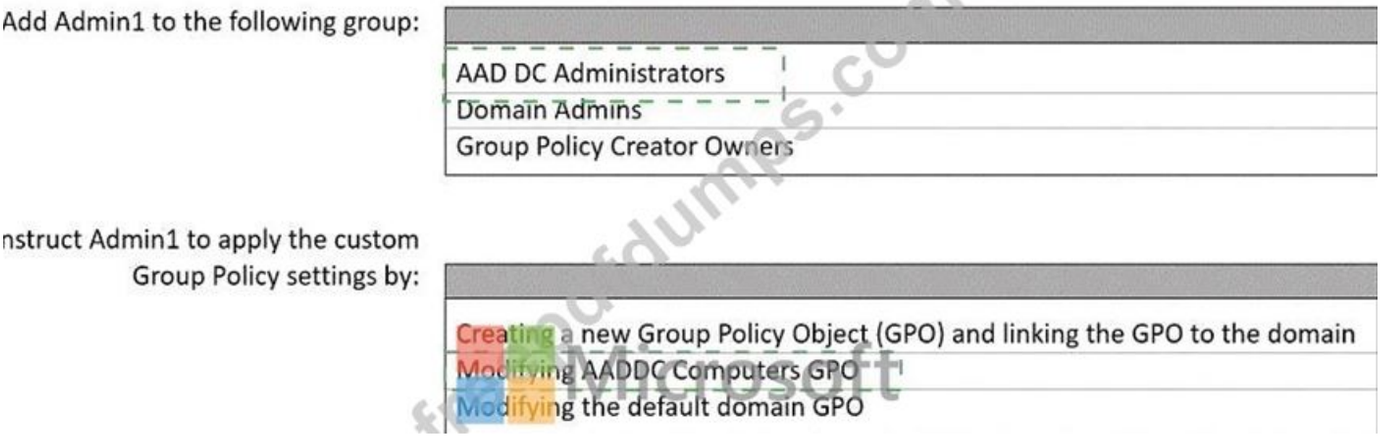

Explanation:

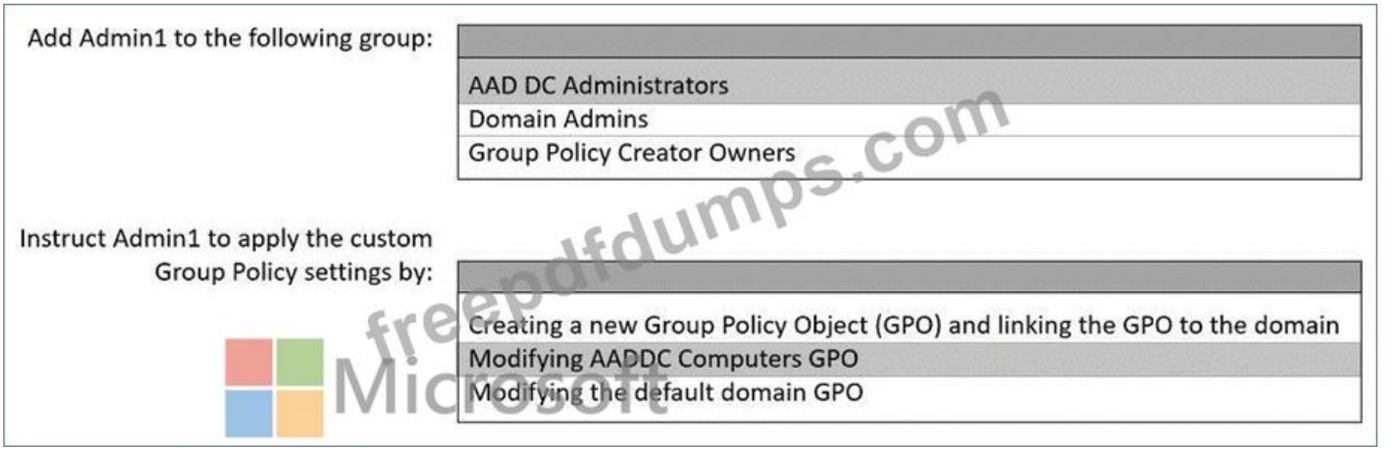

The Administering Windows Server Hybrid Core Infrastructure materials explain that in Azure Active Directory Domain Services (Azure AD DS) you don't get traditional Domain Admins or Group Policy Creator Owners rights. Instead, "administration of the managed domain is delegated to the AAD DC Administrators group. Members of this group can manage Group Policy in the managed domain and administer domain-joined computers." The guide further notes that Azure AD DS automatically creates two built-in OUs and GPOs: "AADDC Computers and AADDC Users, with the corresponding 'AADDC Computers GPO' and 'AADDC Users GPO' already linked." For computer configuration that should apply to all domain-joined machines, the materials state that "you apply or customize computer policy by editing the AADDC Computers GPO (or by creating additional GPOs and linking them to the AADDC Computers OU)." They also emphasize least-privilege and supportability guidance: "Do not modify the Default Domain Policy in Azure AD DS; use the AADDC-scoped GPOs/OUs for managed-domain policy." Putting this together: to let Admin1 deploy custom policy to all computers while honoring least privilege, add Admin1 to AAD DC Administrators (the only group granted GPO management in Azure AD DS) and have them modify the existing 'AADDC Computers GPO' that is already linked to the AADDC Computers OU so the settings flow to every domain-joined computer.

NEW QUESTION: 4

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. Folded and Folder2 only
B. Folder1, volume E, and App1 only
C. Folder1 only
D. Folded. Folder2. App1, and volume E
Answer: D (LEAVE A REPLY)
In the Windows Server hybrid IaaS guidance for AZ-800 (Administering Windows Server Hybrid Core Infrastructure), resizing an Azure VM changes compute characteristics only (vCPU/memory/size family). The learning path on managing Windows Server IaaS VMs states that the VM's OS and any attached data disks persist across stop/deallocate and resize operations, and that drive letters and file data remain intact once the VM starts again. Applications installed on those persistent disks continue to function because the underlying VHDs are unchanged. The only disk that may be transient is the temporary (ephemeral) OS cache or D: temporary disk provided by some sizes; however, managed OS disk (C:) and any managed data disks (such as D: and the newly added E:) retain their contents. Consequently, the folders created on C: and D: remain, the additional data disk that was initialized as volume E remains attached with the same drive letter, and App1 installed on E: is still present after the resize. Therefore, after resizing VM1 you will still have Folder1, Folder2, volume E, and App1, making option D correct.

NEW QUESTION: 5
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 (LEAVE A REPLY)
Reference:
In Administering Windows Server Hybrid Core Infrastructure (AZ-800) guidance on deploying and managing Windows Server IaaS VMs, Microsoft emphasizes automating domain join during provisioning by using Azure VM extensions within an Azure Resource Manager (ARM) template. The supported method for joining Azure VMs to an on-premises AD DS domain is the JsonADDomainExtension (Microsoft.Compute /virtualMachines/extensions). In the ARM template, you provide parameters such as the AD DS domain name, OUPath (optional), User (a domain account with join rights), Password (as a

secure parameter), Restart (true/false), and Options. When the VM is created, the extension executes on first boot and performs the classic AD DS domain join, ensuring every newly deployed VM is joined consistently and at scale.
By contrast, Azure AD Connect (A) only synchronizes identities between AD DS and Azure AD; it does not join Windows Server computers to AD DS. A GPO (B) can configure domain-joined computers after they are in the domain, but it cannot join non-domain systems. Azure management groups (D) provide governance and hierarchy for subscriptions and policies, not machine domain-join operations. Therefore, to ensure that 100 newly deployed Windows Server VMs are automatically joined to the on-premises AD DS domain during deployment, the verified and supported approach is to use an ARM template with the JsonADDomainExtension.

NEW QUESTION: 6
You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Windows Server and has the following disks:
* OSdislcDisk1
o Size: 512 GiB
o Free space: 260 GiB
o Encryption: SSE with PMK
o Storage type: Standard SSD
* Data disk: Disk2
o Size: 512 GiB
o Free space: 45 GiB
o Storage type: Standard HDD
o Encryption: Platform-managed key
You are planning a maintenance strategy for VM1.
You need to identify which task can be performed on Disk2 without causing downtime to VM1.
What should you do on Disk2?
A. Change the encryption type.
B. Decrease the size.
C. Increase the size.
D. Change the storage type to Premium SSD.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 7
You need to meet the technical requirements for VM2.
What should you do?
A. Implement shielded virtual machines.
B. Enable the Guest services integration service.
C. Implement Credential Guard.
D. Enable enhanced session mode.
Answer: (SHOW ANSWER)

In the Administering Windows Server Hybrid Core Infrastructure materials under Hyper-V management, Microsoft specifies that Enhanced Session Mode changes VMConnect from a raw console attach to a connection that uses Remote Desktop Protocol (RDP) to the guest. The guide states that Enhanced Session Mode "uses RDP to establish the VMConnect session so the user must supply credentials for a logon to the guest operating system," and further that it "prevents a second administrator from inheriting an already signed-in console session" because the connection is treated as a new interactive sign-in. In contrast, the default basic console session "attaches directly to the active console without prompting for credentials
," which is exactly the current problem described for VM2.
The same objective area clarifies that other options do not meet the requirement: Guest Services integration only enables file copy and certain host-guest interactions; Credential Guard protects secrets inside Windows by isolating LSASS but does not affect Hyper-V console connection behavior; Shielded VMs provide fabric-level protections and encryption but are not required merely to force credential prompts for VMConnect.
Therefore, to force users to provide credentials when they connect to VM2 and to eliminate inherited console sessions, you should enable Enhanced Session Mode on the Hyper-V host (and ensure the guest supports RDP).

NEW QUESTION: 8
Your network contains an on -premises Active Directory Domain Services (AD DS) domain named contoso.
com The domain contains the objects shown in the following table.

Name	Type
User1	User
Group1	Universal security group
Group2	Domain local security group
Computer1	Computer

You plan to sync contoso.com with an Azure Active Directory (Azure AD) tenant by using Azure AD Connect You need to ensure that all the objects can be used in Conditional Access policies What should you do?
A. Change the scope of Group2 to Universal
B. Clear the Configure device writeback option.
C. Change the scope o' Group1 and Group2 to Global
D. Select the Configure Hybrid Azure AD join option.
Answer: (SHOW ANSWER)
To ensure that all objects, specifically Computer1, can be used in Conditional Access (CA) policies, the environment must support device-based identity in the cloud. In a hybrid scenario, while user objects and security groups (like Group1 and Group2) can be synchronized through standard Azure AD Connect (Microsoft Entra Connect) synchronization cycles, computer objects require specific configuration to become
"identifiable" by Conditional Access.

According to the official study guides for the AZ-800 exam, simply syncing a computer object does not make it a "Hybrid Azure AD joined" device. To enable Computer1 to be used as a target or a condition (e.g.,
"Require Hybrid Azure AD joined device") in a CA policy, you must run the Azure AD Connect wizard and select the Configure Hybrid Azure AD join task. This process configures a Service Connection Point (SCP) in your on-premises Active Directory, which allows Windows 10/11 devices like Computer1 to discover the Azure AD tenant and complete the registration process. Furthermore, while group scopes (Universal vs. Domain Local) are often discussed in sync scenarios, Azure AD Connect by default synchronizes security groups regardless of their scope if they are within the synchronized Organizational Units (OUs). Therefore, the critical step to satisfy the requirement for "all objects"-especially the computer account-is enabling the Hybrid Join feature to establish a cloud-side device identity. This provides the necessary "device signal" that Conditional Access evaluates to grant or deny access.

NEW QUESTION: 9
Your network contains an on-premises Active Directory Domain Services (AD DS) domain named contoso.
com. The domain contains three servers that run Windows Server and have the Hyper-V server role installed.
Each server has a Switch Embedded Teaming (SET) team.
You need to verify that Remote Direct Memory Access (RDMA) and required Windows Server settings are configured properly on each server to support a failover cluster.
What should you use?
A. the validate-DCB cmdlet
B. Server Manager
C. the Get-NetAdapter cmdlet
D. Failover Cluster Manager
Answer: A (LEAVE A REPLY)
In Windows Server environments that use Switch Embedded Teaming (SET) with RDMA/SMB Direct, Microsoft's hybrid server administration guidance specifies validating Data Center Bridging (DCB) and RDMA prerequisites before you place the hosts into a cluster. The Validate-DCB PowerShell cmdlet (from the DataCenterBridging module) is the purpose-built tool to confirm that Priority-based Flow Control (PFC), ETS/QoS policies, and SMB Direct priorities are consistently configured and operational across adapters and vSwitch/SET configurations. The study content emphasizes that resilient RDMA requires correct DCB on every host NIC participating in SMB Direct, and that Validate-DCB "verifies RDMA readiness and flags misconfiguration," including priority mappings and host settings needed for failover clustering. By contrast, Server Manager and Failover Cluster Manager do not test DCB/RDMA health at the NIC policy level, and Get-NetAdapter only reports adapter state (for example, RDMA enabled/disabled) without end-to-end policy validation. Therefore, when you must "verify that RDMA and required Windows Server settings are configured properly on each server" prior to building the cluster, the documented and

supported method is to run Validate-DCB on each Hyper-V host with SET. This aligns with the exam objectives for validating network offloads and SMB Direct for cluster deployments in hybrid infrastructures.

NEW QUESTION: 10

You have a server named Served that runs Windows Server and has the Hyper-V server role installed.

You build Just Enough Administration (JEA) role capabilities and session configuration files.

You need to limit which Hyper-V module cmdlets helpdesk users can use when administering Server1 remotely.

How should you complete the PowerShell command? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Infrastructure documents: =

In Just Enough Administration (JEA), you restrict what remote users can do by creating two artifacts: role capability files (*.psrc) and a session configuration file (*.pssc). The role capability file defines which commands are visible/usable (for example, specific Hyper-V cmdlets via VisibleCmdlets, VisibleFunctions, or VisibleExternalCommands). The session configuration file then maps users or groups to one or more role capabilities using RoleDefinitions.

The Administering Windows Server Hybrid Core Infrastructure material explains that once those files are authored, you must expose a JEA endpoint by registering the session configuration on the target server.

This is accomplished with Register-PSSessionConfiguration, pointing to the PS session configuration file (.

pssc). Registering the endpoint enforces the role capability restrictions when helpdesk users connect (for example, Enter-PSSession -ConfigurationName HyperVJeaHelpDesk -ComputerName Server1).
New-PSSessionConfigurationFile is used earlier to create the .pssc file, but the prompt states the files are already built. Enter-PSSession only starts a session and does not publish the configuration. Therefore, the correct command to limit which Hyper-V module cmdlets helpdesk users can use is:
Register-PSSessionConfiguration -Path .\HyperVJeaConfig.pssc -Name 'HyperVJeaHelpDesk' -Force.

NEW QUESTION: 11
Task 9
You need to replicate a read-only copy of a DNS zone named contosoxom to SRV2.
Answer:
See the solution of this Task below.
Explanation:
Objective:
Create a read-only copy of the DNS zone contoso.com on SRV2.
Step-by-Step Guide: Using a Secondary Zone
Step 1: Log in to SRV2
* Log in to SRV2 (where you want to host the secondary zone) using an account with local administrative privileges.
# Step 2: Open DNS Manager
* Press Windows + R, type dnsmgmt.msc, and press Enter.
# Step 3: Create a Secondary Zone
* In the DNS Manager, expand the server node for SRV2.
* Right-click Forward Lookup Zones and select New Zone.
* The New Zone Wizard opens.
# Step 4: Configure the Secondary Zone
* Zone Type:
* Select Secondary zone and click Next.
* Zone Name:
* Type contoso.com and click Next.
* Master DNS Servers:
* Enter the IP address of the master DNS server that hosts the primary zone (e.g., SRV1's IP).
* Click Next.
* Finish:
* Review the settings and click Finish.
# Step 5: Allow Zone Transfers on the Primary Server
On SRV1 (or the DNS server hosting the primary zone):
* Open DNS Manager.

* Right-click the contoso.com zone and select Properties.
* Go to the Zone Transfers tab.
* Check Allow zone transfers.
* Specify SRV2's IP address (or allow to any server if needed).
# Step 6: Verify Zone Replication
* On SRV2, refresh the Forward Lookup Zones in DNS Manager.
* The contoso.com zone should now appear as a Secondary zone.
* Check the Zone Transfer status to ensure it successfully replicated.

NEW QUESTION: 12
Task 10
You need to configure Hyper-V to ensure that running virtual machines can be moved between SRV1 and SRV2 without downtime.
You do NOT need to move any virtual machines at this time.
Answer:
See the solution o fthis Task below
Explanation:
One possible solution to configure Hyper-V to ensure that running virtual machines can be moved between SRV1 and SRV2 without downtime is to use Live Migration. Live Migration is a feature of Hyper-V that allows you to move a running virtual machine from one host to another without any noticeable interruption of service. To set up Live Migration between SRV1 and SRV2, you need to perform the following steps:
On both SRV1 and SRV2, open Hyper-V Manager from the Administrative Tools menu or by typing virtmgmt.msc in the Run box.
In the left pane, right-click on the name of the server and select Hyper-V Settings.
In the Hyper-V Settings dialog box, select Live Migrations in the navigation pane.
Check the box Enable incoming and outgoing live migrations.
Under Authentication protocol, select the method that you want to use to authenticate the live migration traffic between the servers. You can choose either Kerberos or CredSSP. Kerberos does not require you to sign in to the source server before starting a live migration, but it requires you to configure constrained delegation on the domain controller. CredSSP does not require you to configure constrained delegation, but it requires you to sign in to the source server through a local console session, a Remote Desktop session, or a remote Windows PowerShell session. For more information on how to configure constrained delegation, see Configure constrained delegation.
Under Performance options, select the option that best suits your network configuration and performance requirements. You can choose either TCP/IP or Compression or SMB. TCP/IP uses a single TCP connection for the live migration traffic. Compression uses multiple TCP connections and compresses the live migration traffic to reduce the migration time and network bandwidth usage. SMB uses the Server Message Block (SMB) 3.0 protocol and can leverage

SMB features such as SMB Multichannel and SMB Direct. For more information on how to choose the best performance option, see Choose a live migration performance option.
Under Advanced Features, you can optionally enable the Use any available network for live migration option, which allows Hyper-V to use any available network adapter on the source and destination servers for live migration. If you do not enable this option, you need to specify one or more network adapters to be used for live migration by clicking on the Add button and selecting the network adapter from the list. You can also change the order of preference by using the Move Up and Move Down buttons.
Click OK to apply the settings.
Now, you have configured Hyper-V to enable live migration between SRV1 and SRV2. You can use Hyper-V Manager or Windows PowerShell to initiate a live migration of a running virtual machine from one server to another.

NEW QUESTION: 13
Your company has a main office and 10 branch offices that are connected by using WAN links.
The network contains an Active Directory domain.
All users have laptops and regularly travel between offices.
You plan to implement BranchCache in the branch offices.
In each branch office, you install a server that runs Windows Server and the BranchCache feature. You register the servers in Active Directory.
You need to configure the laptops to use the local BranchCache server automatically. The solution must minimize administrative effort.
Which two Group Policy settings should you configure? To answer, select the settings in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Setting	State
Turn on BranchCache	Not configured
Set BranchCache Distributed Cache mode	Not configured
Set BranchCache Hosted Cache mode	Not configured
Enable Automatic Hosted Cache Discovery by Service Connection...	Not configured
Configure Hosted Cache Servers	Not configured
Configure BranchCache for network files	Not configured
Set percentage of disk space used for client computer cache	Not configured
Set age for segments in the data cache	Not configured
Configure Client BranchCache Version Support	Not configured

Answer:

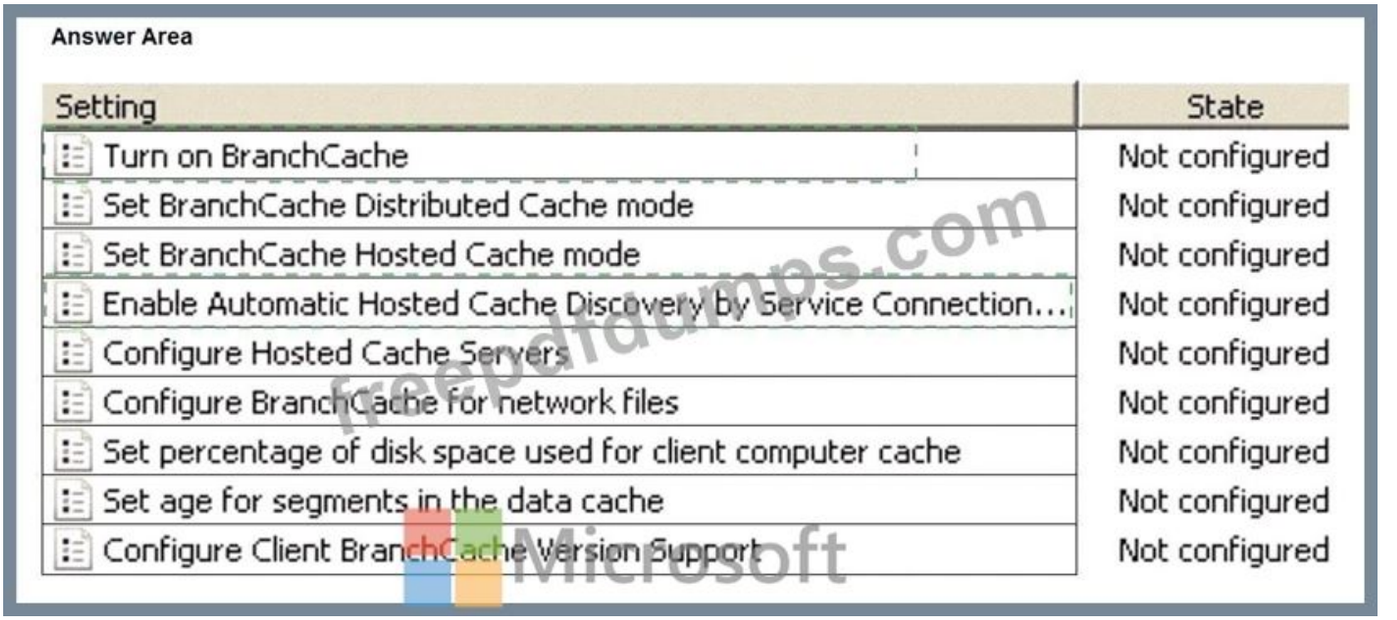

Explanation:

--> Turn on BranchCache

--> Enable Automatic Hosted Cache Discovery by Service Connection ...

NEW QUESTION: 14

Task 5

you need to configure a Group Policy preference to ensure that users in the organizational unit (OU) named Server Admins have a shortcut to a folder named \\srvi.contoso.com\data on their desktop when they sign in to the computers in the domain.

Answer:

See the solution of this Task below.

Explanation:

TASK 5

Objective:

Configure a Group Policy Preference to create a shortcut to \\srvi.contoso.com\data on the desktop of users in the Server Admins OU.

Step-by-Step Guide: Using Group Policy Preferences to Create a Desktop Shortcut

Step 1: Open Group Policy Management Console (GPMC)

Log in to a DC or a management computer with RSAT installed.

Open Group Policy Management (gpmc.msc).

Step 2: Create a New GPO

In the GPMC console, expand the forest and the domain (e.g., contoso.com).

Right-click the OU named Server Admins and select Create a GPO in this domain, and Link it here.

Name the GPO, e.g., Desktop Shortcut for Server Admins.

Step 3: Edit the GPO

Right-click the newly created GPO and select Edit.

This opens the Group Policy Management Editor.

Step 4: Navigate to User Preferences
In the editor, expand:
User Configuration > Preferences > Windows Settings > Shortcuts.
Step 5: Create the Shortcut
Right-click Shortcuts and select New > Shortcut.
In the New Shortcut Properties window:
Action: Create
Name: Data Folder
Target Type: File System Object
Location: Desktop
Target Path: \\srvi.contoso.com\data
Optionally, set an icon or description if you want.
Step 6: Configure Item-Level Targeting (Optional)
If you want to limit this shortcut strictly to specific users/groups, click the Common tab.
Check Item-level targeting and configure conditions (optional).
For this scenario, linking the GPO to the Server Admins OU is usually sufficient.
Step 7: Close and Update
Close the editor.
In GPMC, ensure the GPO is linked to the Server Admins OU.
Force a Group Policy Update on client computers:
On a client computer:
gpupdate /force
Or wait for the next Group Policy refresh cycle.
Step 8: Verify
Log in as a user in the Server Admins OU.
The shortcut to \\srvi.contoso.com\data should appear on the desktop.

NEW QUESTION: 15
Your network contains an Active Directory Domain Services (AD DS) forest. The forest contains the servers shown in the following table.

Name	In domain	Description
Server1	contoso.com	Domain controller, DNS server
Server2	contoso.com	Domain controller, DNS server
Server3	contoso.com	DNS server
Server4	east.contoso.com	Domain controller, DNS server
Server5	east.contoso.com	DNS server

On Server1, you create a DNS zone named Zone1.com as shown in the following exhibit.

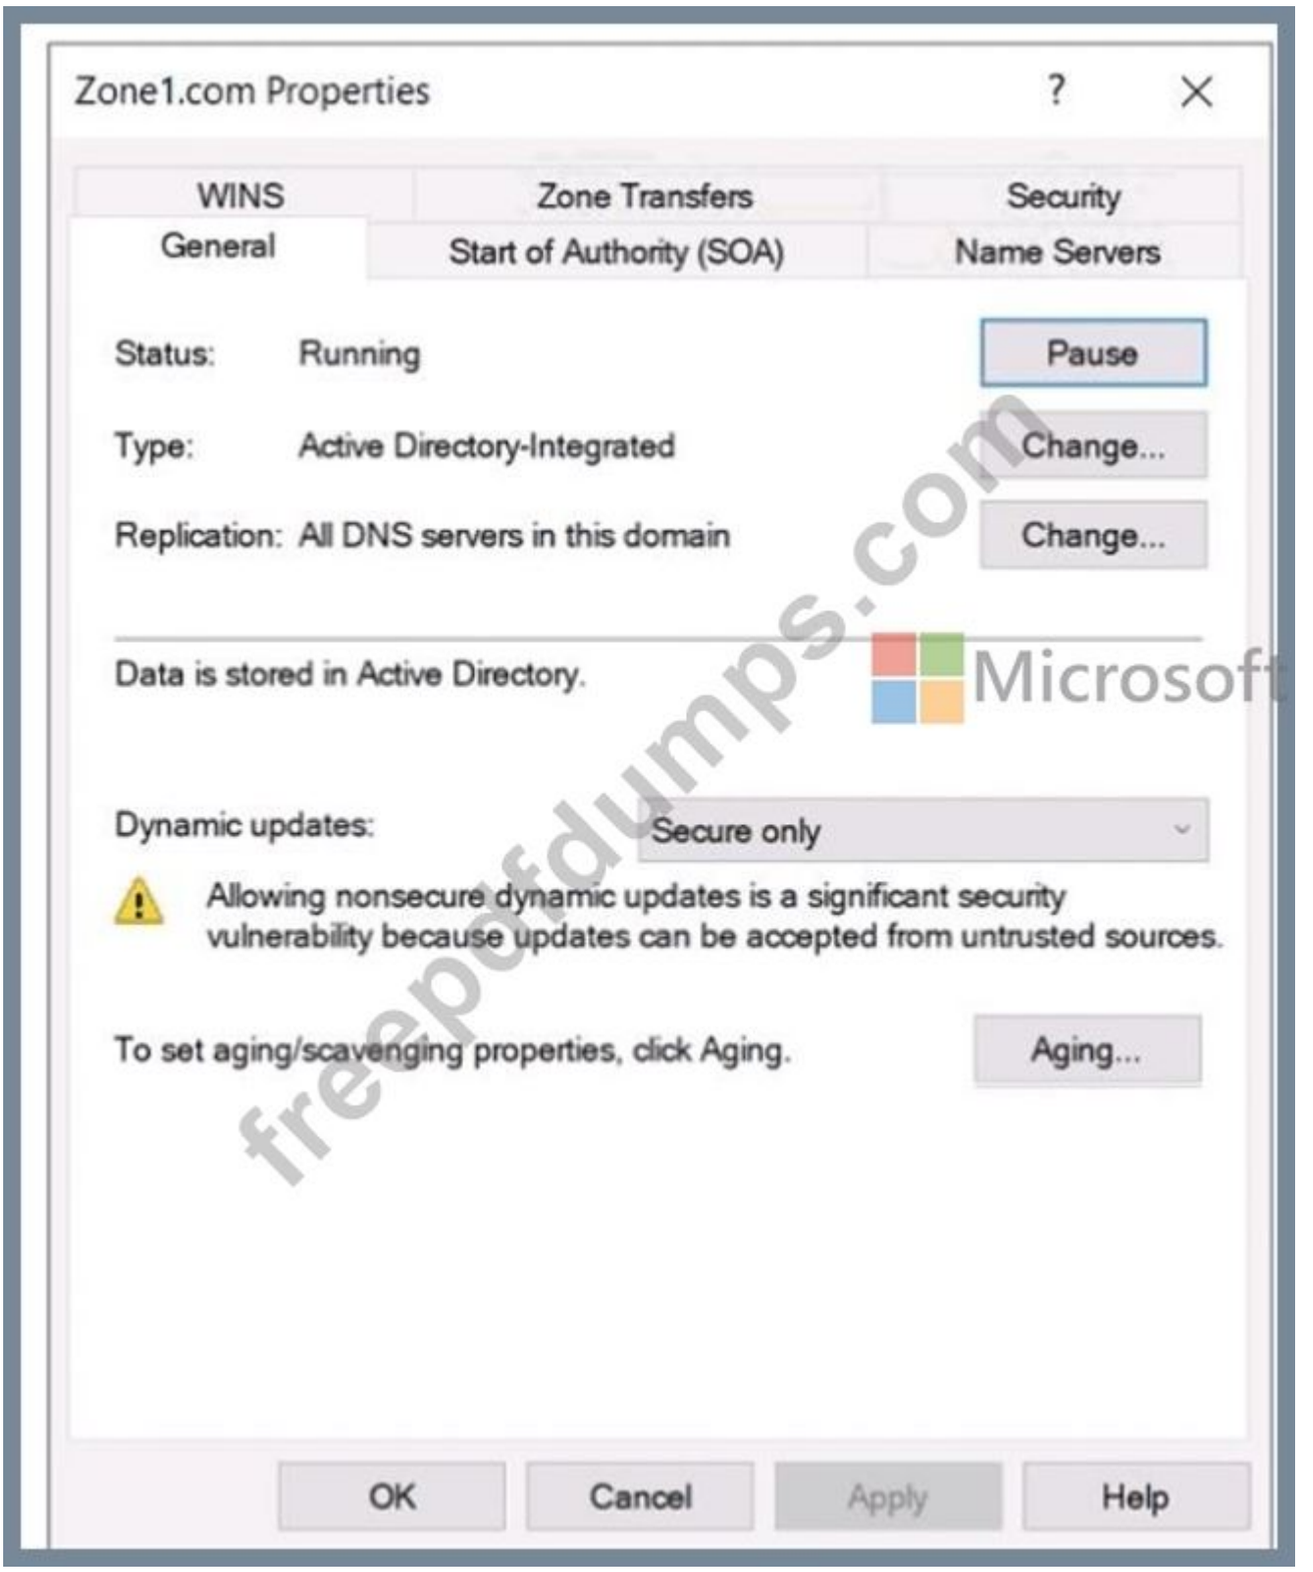

To which DNS servers is Zone1.com replicated?

A. Server2 only

B. Server2 and Server3 only

C. Setver2 and Sen/er4 only

D. Server2. Server3, and Server4 only

E. Server2, Server3, Server4, and Server5

Answer: A (LEAVE A REPLY)

In the Windows Server DNS guidance within Administering Windows Server Hybrid Core Infrastructure, an Active Directory-integrated zone stores its data in AD DS and replicates based on the replication scope selected in the zone's properties. The option "All DNS servers in this domain" replicates the zone to the DomainDNSZones application partition of that domain, which is held only on domain controllers that run the DNS Server role in the same domain. The materials emphasize two key points: (1) AD-integrated zones can be hosted only on DNS servers that are also domain controllers, and (2) replication scope set to "this domain" does not cross to other domains or child domains.
Given the server list: Server1 and Server2 are DCs with DNS in contoso.com; Server3 is DNS-only (not a DC) in contoso.com; Server4/Server5 are in east.contoso.com (a different domain). Therefore, with the scope set to All DNS servers in this domain, replication targets only the other DNS-enabled DCs in contoso.com.
Since you created the zone on Server1, it will replicate to Server2 only; Server3 cannot host AD-integrated zones, and Server4/Server5 are in a different domain and are not in scope.

NEW QUESTION: 16
You have a Windows Server container host named Server1 that has a single disk.
On Server1, you plan to start the containers shown in the following table.

Name	Description
Container1	Container1 is a Windows container that contains a web app in development. The container must **NOT** share a kemel with other containers.
Container2	Container2 is a Linux container that runs a web app. The container requires two static IP addresses.
Container3	Container3 is a Windows container that runs a database. The container requires a static IP address.

Which isolation mode can you use for each container? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

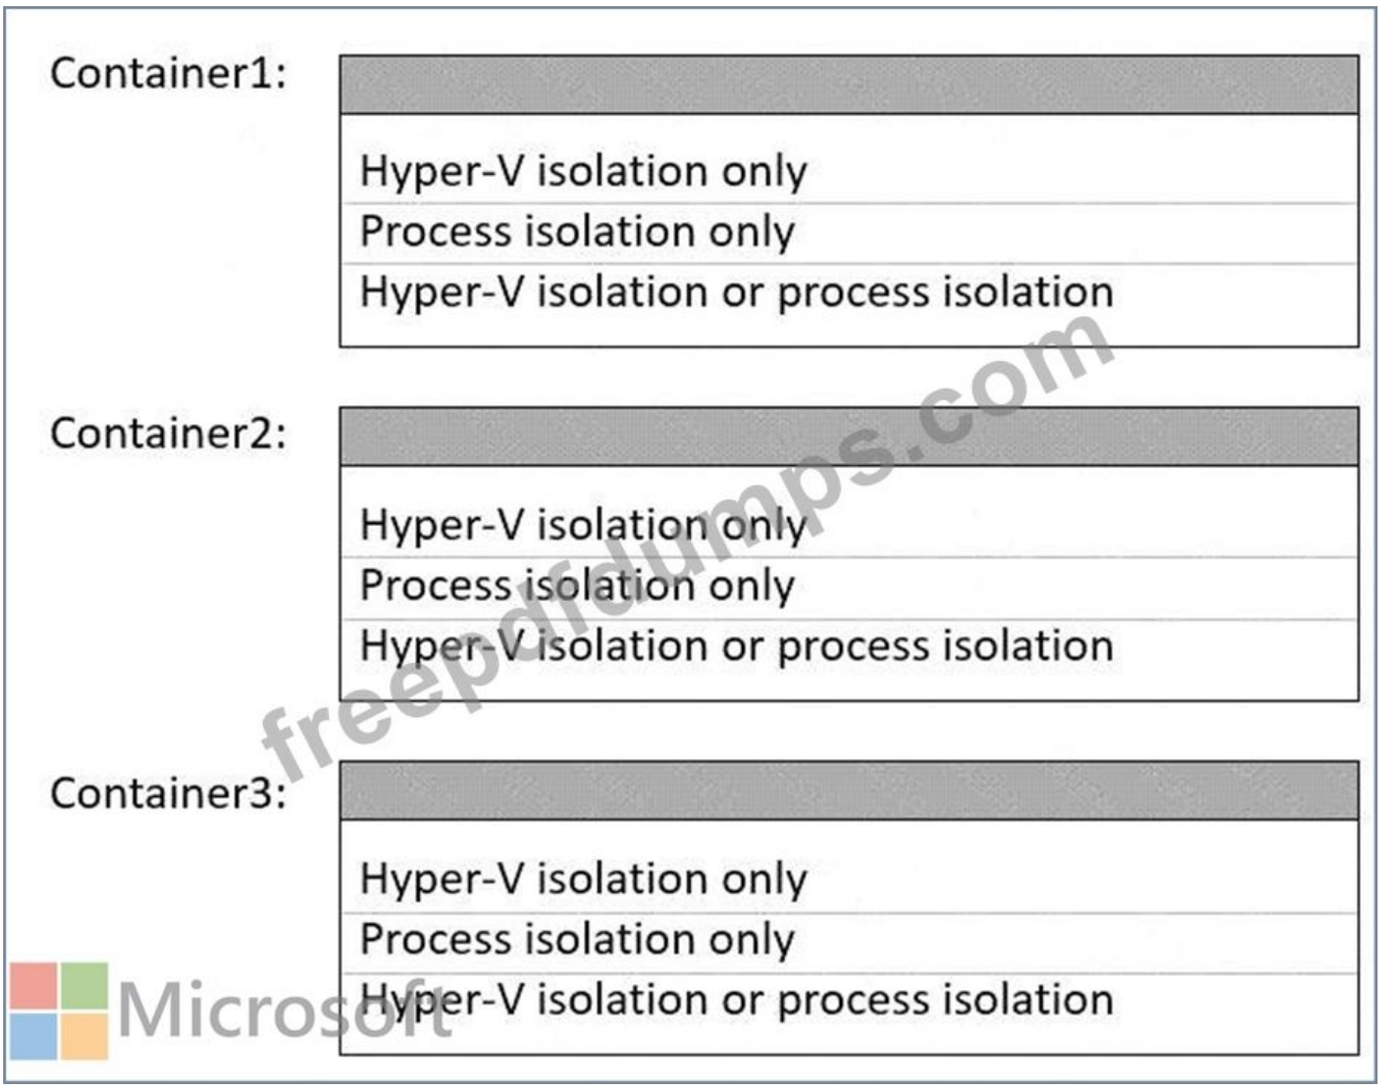

Answer:

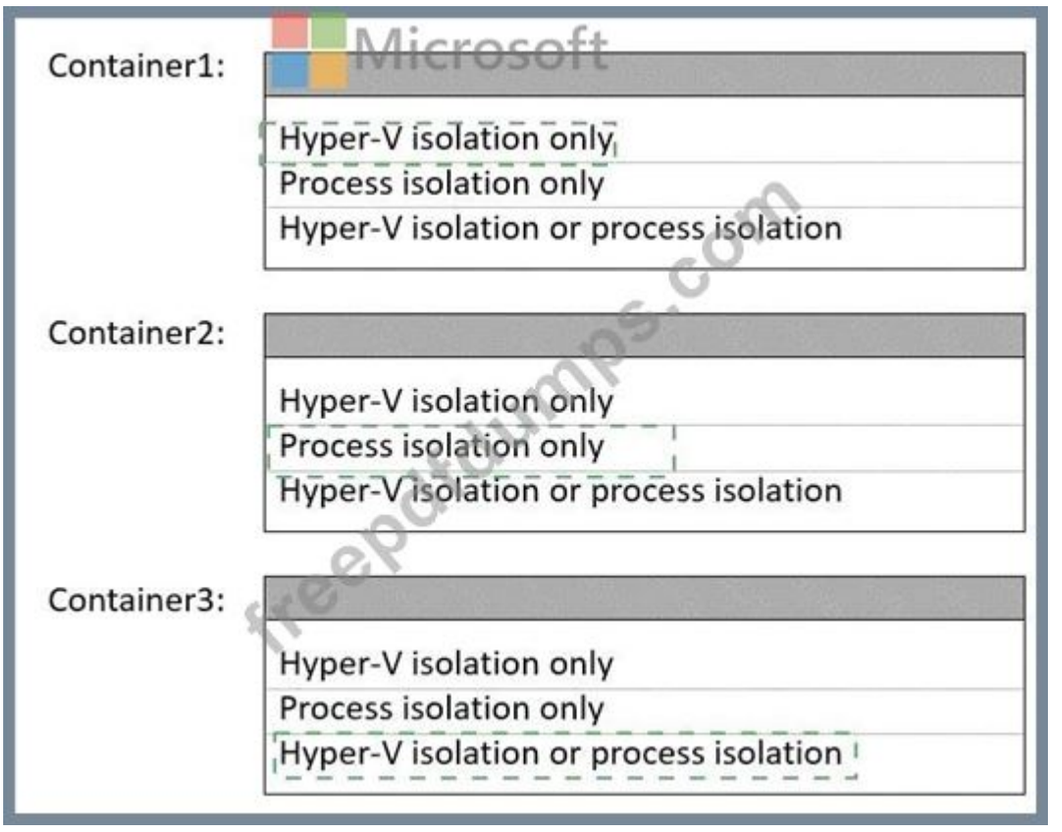

Explanation:

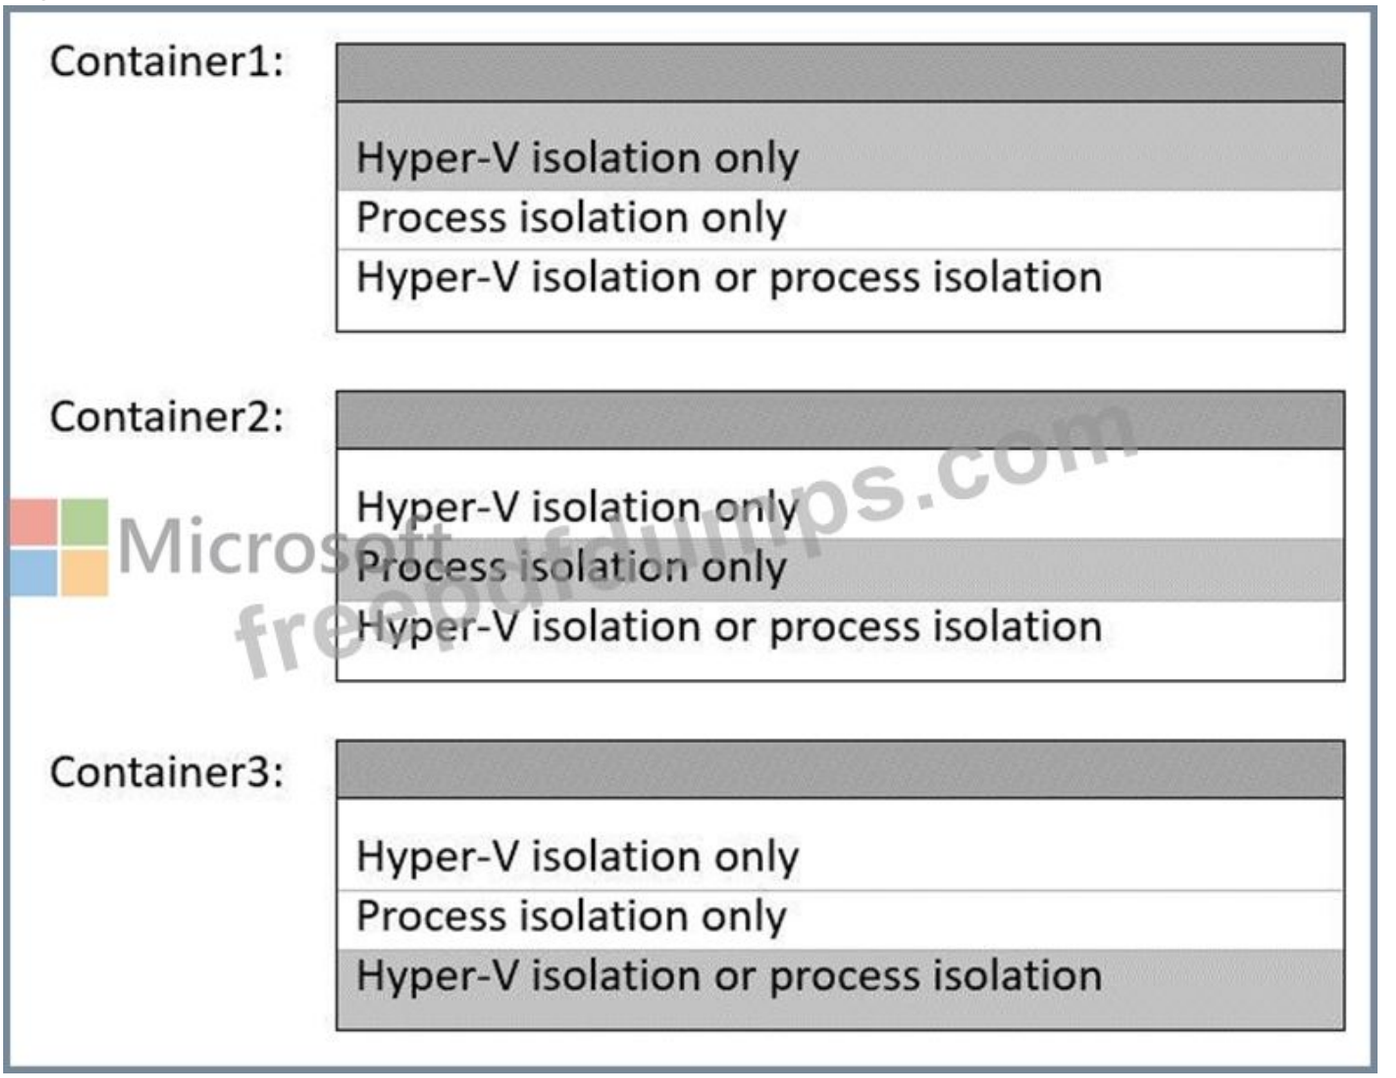

The Administering Windows Server Hybrid Core Infrastructure guidance explains two Windows container isolation modes: process isolation (Windows Server containers) and Hyper-V isolation. With process isolation, containers share the host's kernel, so they run as ordinary processes on the host. With Hyper-V isolation, each container runs inside a lightweight VM with its own kernel; this provides a hard isolation boundary and prevents kernel sharing with other containers. The curriculum further notes that Linux containers on Windows require Hyper-V isolation because they cannot share the Windows kernel.
Applying these rules:
* Container1 "must NOT share a kernel with other containers." The only mode that prevents kernel sharing is Hyper-V isolation, so Container1 must use Hyper-V isolation.
* Container2 is a Linux container. On a Windows Server container host, Linux containers are run via a utility VM and therefore require Hyper-V isolation.
* Container3 is a Windows container that needs a static IP address. The networking modules describe that Windows containers can obtain static IPs when you use supported networks (for example, l2bridge
/transparent) regardless of isolation mode; therefore a Windows workload like a database can run in either process or Hyper-V isolation, depending on the isolation/security you want.
Thus, the correct selections are: Hyper-V isolation only for Container1 and Container2, and Hyper-V isolation or process isolation for Container3.

NEW QUESTION: 17
Your network contains an Active Directory Domain Services (AD DS) domain named adatum.com The domain contains a server named Server1 and the users shown In the following table.

Name	Member of
User1	Group1
User2	Group2
User3	Group3

Server1 contains a folder named D:\Folder1. The advanced security settings for Folder 1 are configured as shown in the Permissions exhibit. (Click the Permissions lab.)

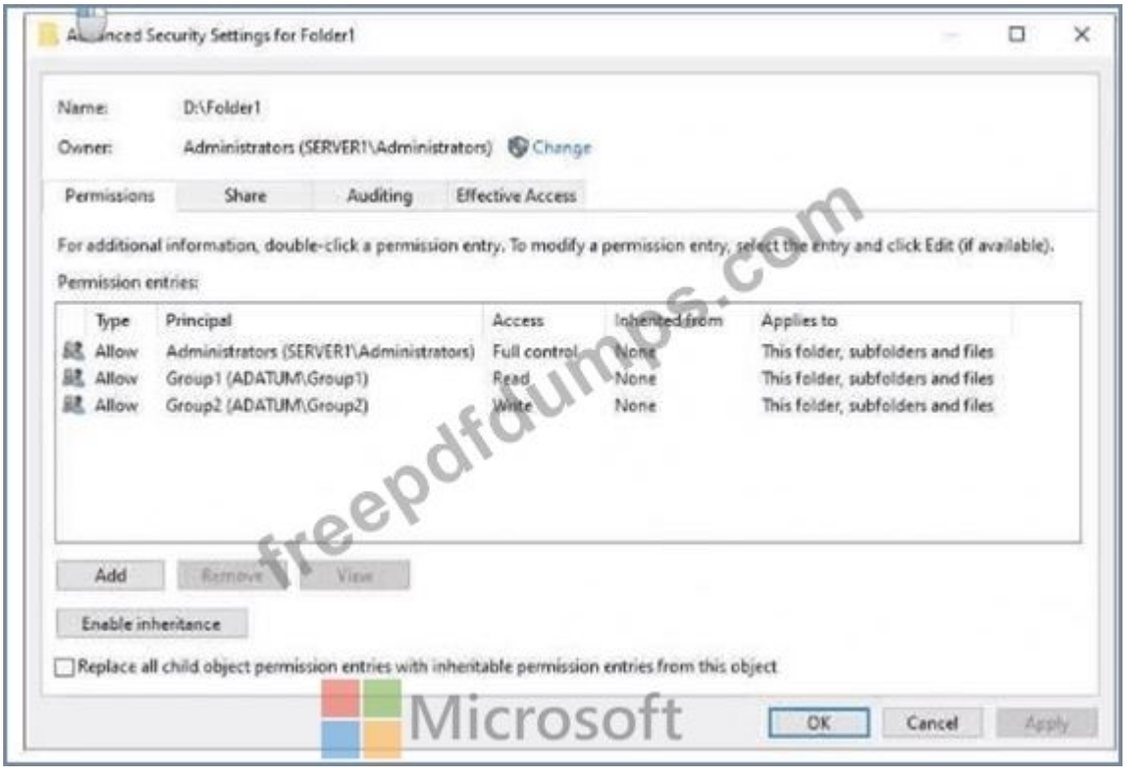

Folder1 is shared by using the following configurations

Group	Permission
Group1	Allow - Change
Group3	Allow - Full Control

```
Path                  : D:\Folder1
Name                  : Share1
ShareType             : FileSystemDirectory
FolderEnumerationMode : Unrestricted
```

Answer Area

Statements	Yes	No
User1 can read the files in Share1.	○	○
User3 can delete files in Share1.	○	○
If User2 connects to \\Server1.adatum.com from File Explorer.	○	○

Answer:

Explanation:

User1 can read the files in Share1. Yes

User3 can delete files in Share1. No

If User2 connects to \Server1.adatum.com from File Explorer, (Share1 will be visible). Yes In Windows file sharing, effective access over SMB = the most restrictive result of Share permissions AND NTFS permissions. The AZ-800 materials emphasize that a user must have sufficient permission on both layers to perform an action. In Folder1, NTFS ACLs grant Group1 =

Read, Group2 = Write, and no entry for Group3. The share "Share1" grants Group1 = Change and Group3 = Full Control.
* User1 (Group1): Share permission "Change" would allow modify over SMB, but NTFS grants only Read. Because the effective permission is the lower of the two, User1 is effectively Read and therefore can read the files.
* User3 (Group3): Although the share grants Full Control, there is no NTFS entry for Group3 (inheritance is disabled), so NTFS denies access. Without NTFS rights (e.g., Modify/Delete), delete is not possible.
* User2 (Group2): NTFS grants Write, but there is no share permission for Group2, so User2 cannot access content through the share. However, the share's FolderEnumerationMode = Unrestricted (i.e., Access-Based Enumeration is off). As covered in the hybrid core guide, when ABE is disabled, users can see items even if they lack permissions. Thus, when User2 browses \Server1.adatum.com, Share1 is visible (opening it will result in Access Denied).
Therefore: Yes (User1 read), No (User3 delete), Yes (User2 sees Share1).

NEW QUESTION: 18
Your network contains an Active Directory Domain Services (AD DS) domain. The domain contains two servers named Server1 and Server2 and the users shown in the following table.

Name	Member of
User1	Contoso\Administrators
User2	Contoso\Remote Management Users
User3	Server2\Administrators
User4	Server2\Remote Management Users

Which users can establish a PowerShell remoting session from Server1 to Server2?
A. User1 and User3 only
B. User2 and User4 only
C. User3 and User4 only
D. User1, User3. and User4 only
E. User1, User2, User3, and User4
Answer: (SHOW ANSWER)
The remoting prerequisites in the AZ-800 curriculum state that WinRM/PowerShell remoting to a target computer is permitted for local Administrators and Remote Management Users on the target. The documentation notes: "Users who are members of the local Administrators group or Remote Management Users group on the destination can establish PowerShell remoting sessions." In a domain, high-privilege domain administrative groups are, by default, granted local administrator rights on domain-joined servers.
Applying this: User3 is in Server2\Administrators (local admin) # allowed. User4 is in Server2\Remote Management Users # allowed. User1 belongs to the domain Administrators group, which confers administrator privileges on domain-joined servers, enabling remoting to Server2. User2, however, is only in the domain "Remote Management Users" group, not the local group on Server2; domain membership alone does not grant the required local right. Therefore,

the users who can open a PowerShell remoting session from Server1 to Server2 are User1, User3, and User4.

NEW QUESTION: 19

Your network contains the segments shown in the following table.

Name	IPv4 address space	Gateway
Segment1	172.16.1.0/24	172.16.1.1
Segment2	172.16.2.0/24	172.16.2.1

You have servers that run Windows Server and are configured as shown in the following table.

Name	IPv4 address	Connected to	Windows Defender Firewall configuration
Server1	172.16.1.10	Segment1	Allow ICMP traffic
Server2	172.16.2.2	Segment1	Allow ICMP traffic
Server3	172.16.2.20	Segment2	Allow ICMP traffic

You deploy a server named Server4 that runs Windows Server and has a static IP address of 172.16.1.1. You connect Server4 to Segment1.

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

NOTE: Each correct selection is worth one point.

Answer:

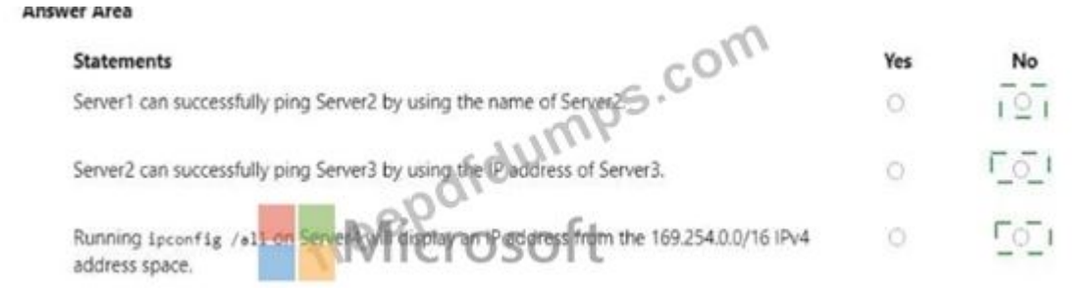

Explanation:

< Server1 can successfully ping Server2 by using the name of Server2. No Server2 can successfully ping Server3 by using the IP address of Server3. No Running ipconfig /all on Server4 will display an IP address from the 169.254.0.0/16 IPv4 address space.

No

The Windows Server networking objectives explain that hosts determine whether a destination is local by comparing the destination IP to their own subnet and mask. If the destination is outside the local subnet, traffic is sent to the default gateway; if it is within the local subnet, the host uses ARP to resolve the destination MAC on the local segment. In this scenario, Server1 (172.16.1.10/24) and Server2 (172.16.2.2/24) are on different subnets. Even though Server2 is physically attached to Segment1, its IP places it in the

172.16.2.0/24 network. When Server1 pings Server2 (by name or IP), name resolution returns 172.16.2.2, which Server1 treats as remote and forwards to the gateway 172.16.1.1. The router will forward toward Segment2 (172.16.2.0/24), where Server2 is not physically present, so communication fails. Similarly, Server2 # Server3: Server2 believes 172.16.2.20 is local and issues an ARP on Segment1; ARP cannot cross to Segment2, so no reply-ping fails. The materials also note that APIPA (169.254.0.0/16) is assigned only when DHCP fails on DHCP-configured interfaces. Because Server4 was configured with a static address (172.16.1.1), ipconfig /all will show that static IP, not an APIPA address. Additionally, assigning 172.16.1.1 to Server4 conflicts with the listed gateway, further preventing routed communication but does not change the fact that Server4 uses its configured static IP.

NEW QUESTION: 20
You need to ensure that access to storage1 for the Marketing OU users meets the technical requirements.
What should you implement?
A. Microsoft Entra Connect cloud sync
B. Active Directory Federation Services (AD FS)
C. Microsoft Entra Connect in staging mode
D. Microsoft Entra Connect in active mode
Answer: (SHOW ANSWER)
The Administering Windows Server Hybrid Core Infrastructure content notes that Microsoft Entra Connect cloud sync uses lightweight agents and supports synchronizing specific OUs from multiple forests into a single tenant. The guide highlights that cloud sync "is designed for multi-forest or cross-organization scenarios and allows scoping to selected OUs and groups," enabling granular onboarding of just the identities you need. The requirement says "the users in the Marketing OU (in Fabrikam's forest) must have access to storage1." Granting access to Azure Files using Entra-based authorization requires that those users exist in the same Entra tenant. Cloud sync enables importing only the Fabrikam Marketing OU into A. Datum's tenant without establishing full trust for all users. Azure AD Connect in active or staging mode would target the A.
Datum forest and isn't intended for selectively bringing in a separate partner forest OU; AD FS is a federation solution and does not create tenant objects for ACLs on Azure resources like Azure Files. Therefore, implement Microsoft Entra Connect cloud sync and scope it to the Marketing OU to meet the requirement.

NEW QUESTION: 21
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 are planning the deployment of DNS to a new network.
You have three internal DNS servers as shown in the following table.

Name	Location	IP address	Local DNS zone
Server1	Montreal	10.0.1.10	contoso.local
Server2	Toronto	10.0.2.10	east.contoso.local
Server3	Seattle	10.0.3.10	west.contoso.local

The contoso.local zone contains zone delegations for east.contoso.local and west.contoso.local.
All the DNS servers use root hints.
You need to ensure that all the DNS servers can resolve the names of all the internal namespaces and internet hosts.
Solution: You configure Server2 and Server3 to forward DNS requests to 10.0.1.10.
Does this meet the goal?
A. Yes
B. No
Answer: A (LEAVE A REPLY)
In the Windows Server DNS planning guidance from Administering Windows Server Hybrid Core Infrastructure, forwarders can be used to centralize name resolution through a designated DNS server. The guide states that a DNS server can be configured to "forward queries it cannot resolve to one or more upstream DNS servers" and that this is often used to "centralize Internet and internal namespace resolution through an authoritative hub server." In this scenario, Server1 hosts the contoso.local parent zone and already contains delegations to east.contoso.local and west.contoso.local. If Server2 and Server3 are set to forward unresolved queries to 10.0.1.10 (Server1), they will resolve:
* Internal names-Server1 is authoritative for the parent and, through delegations, can refer requests to the appropriate child-zone servers.
* Internet names-Server1 uses root hints and can resolve external hosts, with Server2/Server3 receiving the answers via forwarding.
The study materials emphasize: "Delegations enable a parent zone to direct queries to child zones," and
"forwarders do not break authority-authoritative data is answered locally; only unresolved names are forwarded." Thus, configuring Server2 and Server3 to forward to Server1 satisfies the requirement that all DNS servers resolve all internal namespaces and Internet hosts, while keeping the design simple and consistent with DNS best practices.

NEW QUESTION: 22
Your network contains an Active Directory Domain Services (AD DS) domain named contoso.com.
The network contains the servers shown in the following table.

Name	Role	Domain/workgroup	Operating system
DC1	Active Directory Domain Services, DNS Server	contoso.com	Windows Server
Server1	DHCP Server	contoso.com	Windows Server
Server2	None	contoso.com	Windows Server Core
Server3	None	Workgroup1	Windows Server Core

You plan to implement IP Address Management (IPAM).

You need to use the Group Policy based provisioning method for managed servers. The solution must support server discovery.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

The Administering Windows Server Hybrid Core Infrastructure materials state the placement and provisioning requirements for IP Address Management (IPAM). Specifically:

* "The IPAM server must be installed on a domain member computer. You cannot install the IPAM server feature on a domain controller."
* "IPAM manages and discovers domain controllers, DHCP servers, and DNS servers in the domain.

Servers in a workgroup are not supported as managed servers."

* "When you choose Group Policy-based provisioning, IPAM creates and links GPOs that configure the required settings on managed DC/DHCP/DNS servers so that IPAM can perform inventory, event collection, and address space management." Applying these rules:

* Server2 is a domain-joined Windows Server Core member with no conflicting roles, satisfying the guidance to avoid installing IPAM on a DC and aligning with the recommendation to place IPAM on a dedicated member server. DC1 is a domain controller (and DNS), so it must not host

IPAM. Server1 (DHCP) could host IPAM but best practice is to use a dedicated server. Server3 is in a workgroup, so it cannot host IPAM or be managed by it.
* For GPO-based provisioning and server discovery, the managed servers that must be provisioned are the infrastructure role holders: DC1 (AD DS/DNS) and Server1 (DHCP). These are exactly the servers IPAM discovers and manages via the created GPOs.

NEW QUESTION: 23
You have an on-premises Active Directory Domain Services (AD DS) domain that syncs with an Azure Active Directory (Azure AD) tenant You have an on-premises web app named WebApp1 that only supports Kerberos authentication.
You need to ensure that users can access WebApp1 by using their Azure AD account. The solution must minimize administrative effort.
What should you configure? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

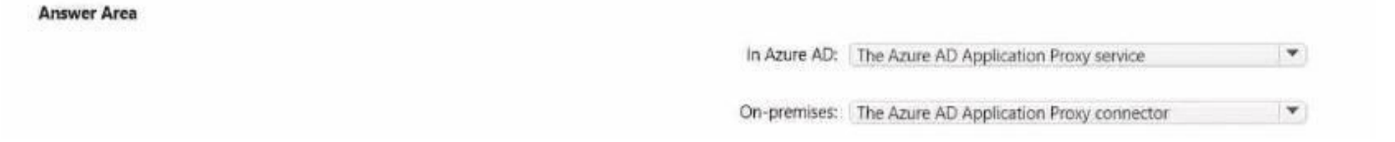

In the Administering Windows Server Hybrid Core Infrastructure objectives for publishing on-premises apps to cloud identities, Azure AD Application Proxy is the prescribed solution for exposing internal, Kerberos- based web apps to Azure AD users with minimal changes. The guidance explains that Application Proxy is a cloud service in Azure AD which publishes internal HTTP/HTTPS applications and supports Kerberos Constrained Delegation (KCD) so that "users authenticate with Azure AD and the service performs single sign-on to the on-premises app using Kerberos." To enable it, you turn on the Application Proxy service in the Azure AD tenant and configure the application object for passthrough/Pre-Auth and KCD. The on-premises requirement is to install the Azure AD Application Proxy connector on a domain-joined Windows Server that can reach the internal application; the connector establishes outbound-only connections to the

service, avoiding inbound firewall changes and minimizing administrative effort. The materials also contrast this with Web Application Proxy (WAP) for AD FS publishing, which requires additional infrastructure and does not leverage Azure AD cloud pre-authentication. Therefore, to let Azure AD accounts reach an on-premises Kerberos-only app with least effort, configure the Azure AD Application Proxy service (in Azure AD) and deploy the Azure AD Application Proxy connector (on-premises), then enable KCD for the published app.

NEW QUESTION: 24
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.
Your network contains an Active Directory Domain Services (AD DS) forest. The forest contains three Active Directory sites named Site1, Site2, and Site3. Each site contains two domain controllers. The sites are connected by using DEFAULTIPSITELINK.
You open a new branch office that contains only client computers.
You need to ensure that the client computers in the new office are primarily authenticated by the domain controllers in Site1.
Solution: You create an organization unit (OU) that contains the client computers in the branch office. You configure the Try Next Closest Site Group Policy Object (GPO) setting in a GPO that is linked to the new OU.
Does this meet the goal?
A. Yes
B. No
Answer: B (LEAVE A REPLY)
The DC Locator behavior covered in the Active Directory sites and services section of Administering WSHCI explains that a client's logon target is determined first by its site membership, which is computed from the IP subnet-to-site mapping in AD DS. Only if no domain controller is available in the client's own site does DC Locator consider other sites based on site-link costs ("closest site"). The policy Try Next Closest Site merely instructs clients to fall back to the nearest site by cost when they have no local DC; it does not pin clients to a specific site. Linking that GPO to an OU that contains the branch computers does not change their computed site, because site assignment is not influenced by OU/GPO scope-it is solely driven by subnet objects. As the guide notes: "Client site affinity is determined by IP subnet mapping; Group Policy scope doesn't affect site discovery." Since the requirement is to ensure branch clients are primarily authenticated by Site1, you must make those clients members of Site1 via subnet mapping (or adjust site link costs accordingly). Therefore, enabling Try Next Closest Site on an OU does not meet the goal.

NEW QUESTION: 25

You have an on-premises server named Server! that runs Windows Server.
You have an Azure subscription that contains a virtual network named VNet1.
You need to connect Server1 to VNet1 by using Azure Network Adapter.
What should you use?
A. Azure AD Connect
B. Device Manager
C. the Azure portal
D. Windows Admin Center
Answer: D (LEAVE A REPLY)
The Administering Windows Server Hybrid Core Infrastructure study guides describe Azure Network Adapter (ANA) as "a Windows Admin Center (WAC) workflow that creates a point-to-site VPN from an on- premises Windows Server to an Azure virtual network." The guidance emphasizes: "ANA is initiated and configured exclusively from Windows Admin Center; it automates creation of the Azure VPN gateway side and the client configuration on the server." The Azure portal does not expose an "ANA" button for a standalone server, Device Manager has no role, and Azure AD Connect is for directory synchronization, not networking. The steps include registering WAC with Azure, selecting Add > Azure Network Adapter on the target server, choosing the subscription/resource group/VNet (e.g., VNet1), and letting WAC provision the P2S configuration and Windows VPN client profile. Therefore, to connect Server1 to VNet1 using Azure Network Adapter, you must use Windows Admin Center.

NEW QUESTION: 26
You have a server named Server1 that runs Windows Server. The disks on Server1 are configured as shown in the following exhibit.

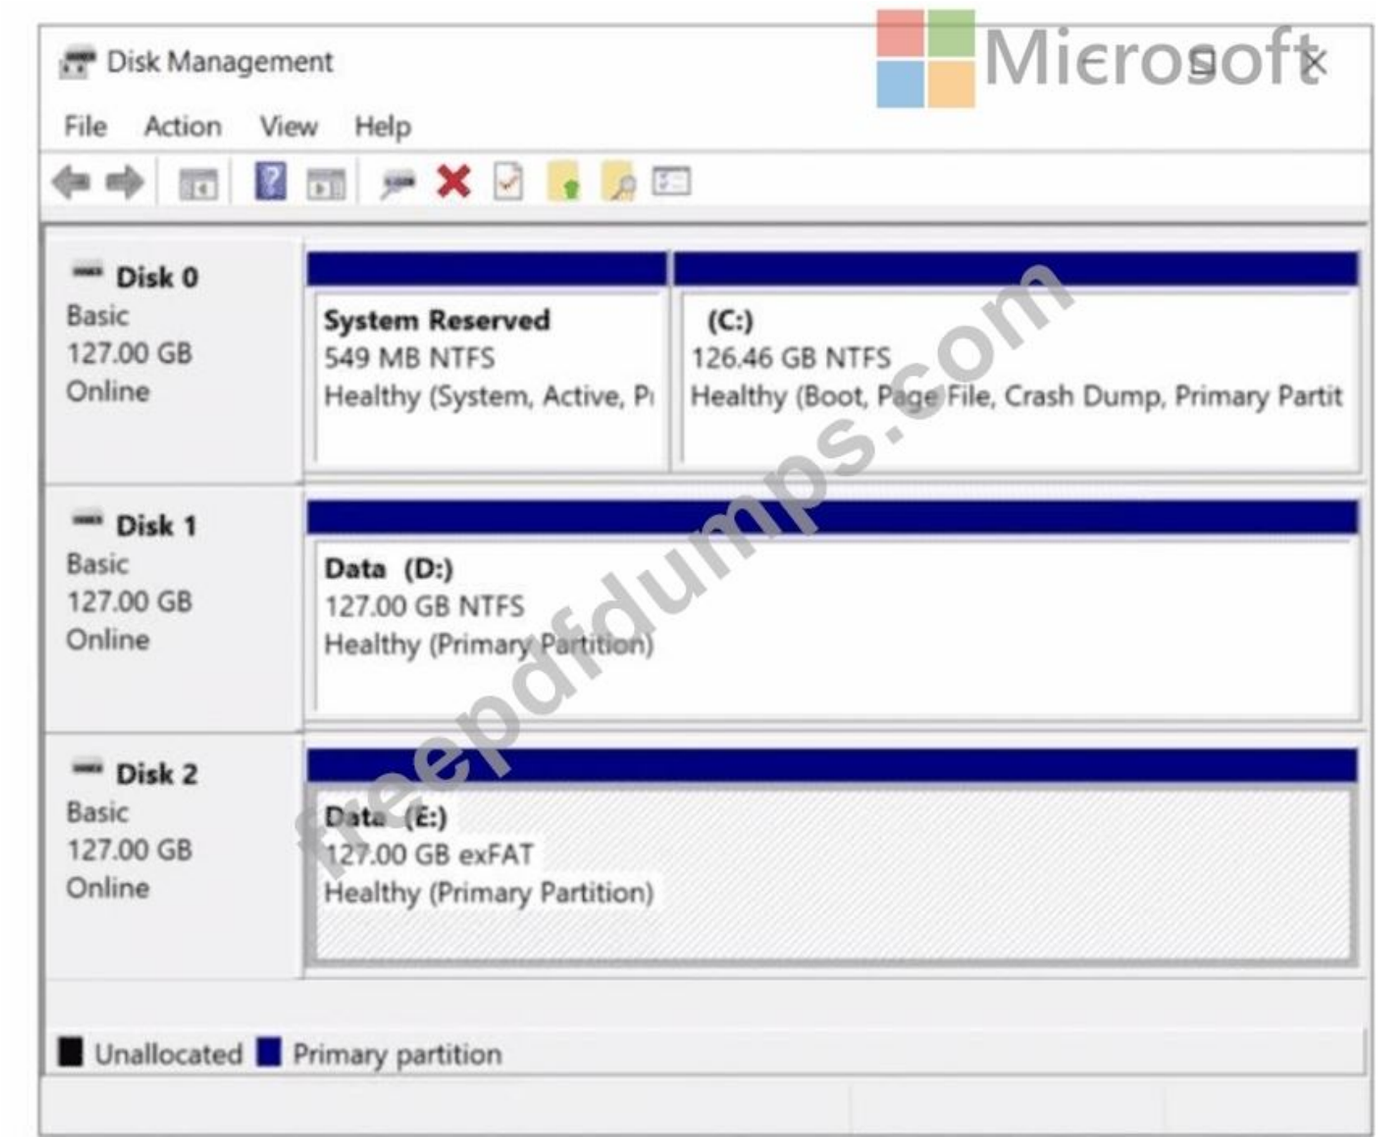

You need to convert volume E to ReFS. The solution must meet the following requirements:

* Preserve the existing data on volume E.
* Minimize administrative effort.

What should you do first?

A. Take Disk 2 offline.

B. Back up the data on volume E.

C. Convert Disk 2 to a dynamic disk.

D. Runconvert.exe.

Answer: (SHOW ANSWER)

The Windows Server storage objectives in Administering Windows Server Hybrid Core Infrastructure make clear that ReFS cannot be converted in-place from FAT/FAT32/exFAT. The guide notes that the only supported in-place file-system conversion utility is convert.exe for FAT # NTFS, and "there is no tool to convert an existing volume to ReFS without reformatting." The prescribed approach is: "Back up the data, recreate the volume with the ReFS file system, and then restore the data." Because the screenshot shows Disk

2, Volume E: formatted as exFAT, preserving data while moving to ReFS first requires a backup; afterwards you would reformat E: as ReFS and restore. Taking the disk offline or converting the

disk to dynamic does not change the file system. Running convert.exe is irrelevant because it only targets FAT to NTFS, not ReFS.
Thus, to meet the requirements (preserve data, minimize effort), the first step is to back up the data on E:, then reformat to ReFS and restore.

NEW QUESTION: 27
You have an on-premises server named Server 1 that runs Windows Server. You have an Azure subscription that contains a virtual network named VNetl. You need to connect Server1 to VNetl by using Azure Network Adapter. What should you use?
A. the Azure portal
B. Azure AD Connect
C. Device Manager
D. Windows Admin Center
Answer: D (LEAVE A REPLY)
The hybrid networking module in AZ-800 details Azure Network Adapter (ANA) as a Windows Admin Center (WAC) workflow that creates a point-to-site (P2S) VPN from an on-premises Windows Server to an Azure virtual network. The documentation states that ANA is initiated directly from Windows Admin Center
, which automates provisioning of the necessary Azure resources (e.g., VPN gateway configuration and certificates) and installs/configures the VPN client on the server, establishing secure connectivity to the target VNet. The Azure portal does not install or configure the on-premises VPN client for a standalone server; Azure AD Connect is used for identity synchronization and is unrelated to network tunneling; and Device Manager is not used for configuring VPN or Azure connectivity. In practice, you add the server to WAC, open the Network tool, and select Add Azure Network Adapter, sign in to Azure, choose VNet1, and complete the wizard. This is the prescribed, streamlined method for connecting Server1 to VNet1 using Azure Network Adapter, validating Windows Admin Center as the correct choice.

NEW QUESTION: 28
You have a server named Server 1 that runs Windows Server and has the Hyper-V server role installed.
Server1 hosts a virtual machine named VM1. Server1 has an NV Me storage device that is assigned to VM1 by using Discrete Device Assignment.
You need to make the device available to the host.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

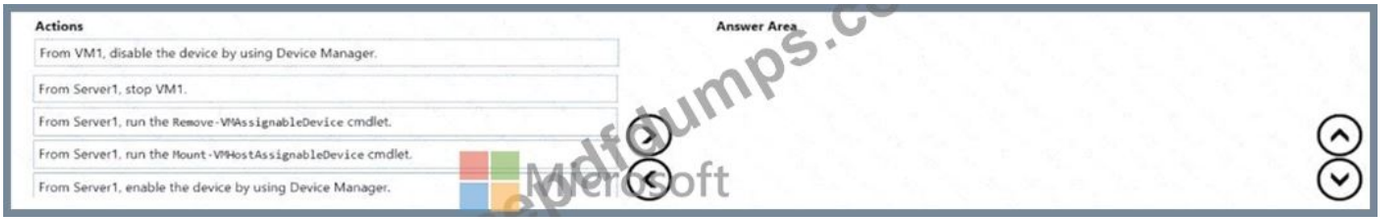

Answer:

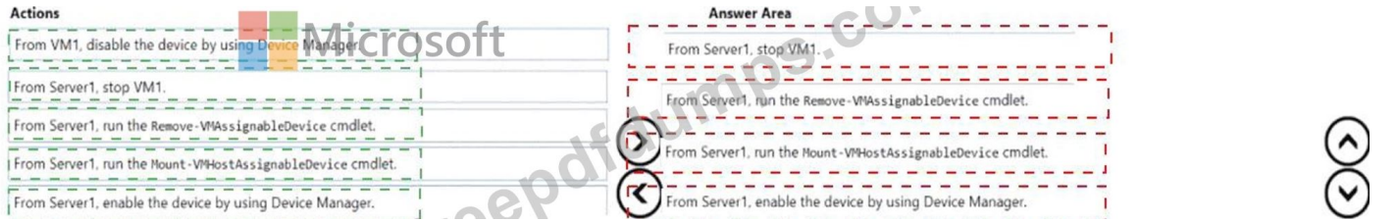

Explanation:

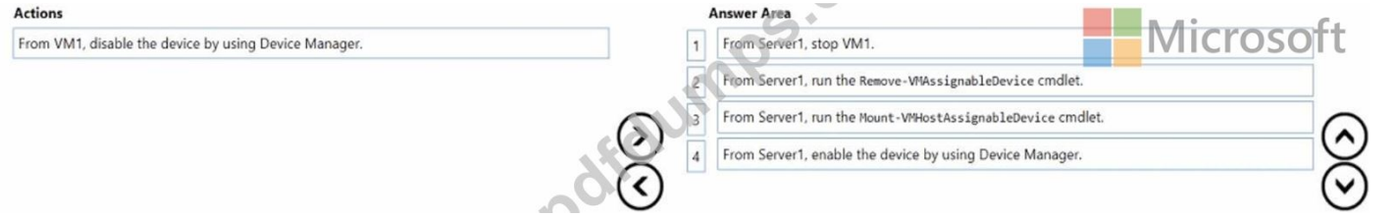

The Administering Windows Server Hybrid Core Infrastructure guidance for Hyper-V Discrete Device Assignment (DDA) explains that a PCIe/NVMe device passed through to a VM becomes detached from the host PnP stack and "owned" by the guest. To reclaim the device for the host you must first turn off the VM because "changes to DDA device assignments require the virtual machine to be in an Off state." After the VM is stopped, you remove the association using Remove-VMAssignableDevice (the inverse of assigning with Add-VMAssignableDevice). The device is then returned from the VM's configuration. Next, you return the device to the host by calling Mount-VMHostAssignableDevice, which "re-attaches the PCI device to the host and makes it available to the host operating system." Finally, because the host previously dismounted the device to make it assignable, Windows on the host requires a PnP re-enable: use Device Manager (or a rescan) to enable the hardware so the host can use it again.
Therefore, the correct sequence to make the NVMe device available to Server1 is: stop VM1, run Remove- VMAssignableDevice, run Mount-VMHostAssignableDevice, and then enable the device in Device Manager on the host.

NEW QUESTION: 29

Task 7

You need to monitor the security configuration of DC1 by using Microsoft Defender for Cloud.
The required source files are located in a folder named \\dc1.contoso.com\install.

Answer:

See the solution of this Task below.

Explanation:

One possible solution to monitor the security configuration of DC1 by using Microsoft Defender for Cloud is to use the Guest Configuration feature. Guest Configuration is a service that audits settings inside Linux and Windows virtual machines (VMs) to assess their compliance with your organization's security policies. You can use Guest Configuration to monitor the security baseline settings for Windows Server in the Microsoft Defender for Cloud portal by following these steps:
On DC1, open a web browser and go to the folder named \dc1.contoso.com\install. Download the Guest Configuration extension file (GuestConfiguration.msi) and save it to a local folder, such as C:\Temp.

Run the Guest Configuration extension file and follow the installation wizard. You can choose to install the extension for all users or only for the current user. For more information on how to install the Guest Configuration extension, see Install the Guest Configuration extension.
After the installation is complete, sign in to the Microsoft Defender for Cloud portal (2).
In the left pane, select Security Center and then Recommendations.
In the recommendations list, find and select Vulnerabilities in security configuration on your Windows machines should be remediated (powered by Guest Configuration).
In the Remediate Security Configurations page, you can see the compliance status of your Windows VMs, including DC1, based on the Azure Compute Benchmark. The Azure Compute Benchmark is a set of rules that define the desired configuration state of your VMs. You can also see the number of failed, passed, and skipped rules for each VM. For more information on the Azure Compute Benchmark, see Microsoft cloud security benchmark: Azure compute benchmark is now available.
To view the details of the security configuration of DC1, click on the VM name and then select View details.
You can see the list of rules that apply to DC1 and their compliance status. You can also see the severity, description, and remediation steps for each rule. For example, you can see if DC1 has the latest security updates installed, if the firewall is enabled, if the password policy is enforced, and so on.
To monitor the security configuration of DC1 over time, you can use the Compliance over time chart, which shows the trend of compliance status for DC1 in the past 30 days. You can also use the Compliance breakdown chart, which shows the distribution of compliance status for DC1 by rule severity.
By using Guest Configuration, you can monitor the security configuration of DC1 by using Microsoft Defender for Cloud and ensure that it meets your organization's security standards. You can also use Guest Configuration to monitor the security configuration of other Windows and Linux VMs in your Azure environment.

NEW QUESTION: 30
Your network contains two VLANs for client computers and one VLAN for a datacenter Each VLAN is assigned an IPv4 subnet Currently, all the client computers use static IP addresses.
You plan to deploy a DHCP server to the VLAN in the datacenter.
You need to use the DHCP server to provide IP configurations to all the client computers.
What is the minimum number of scopes and DHCP relays you should create? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

* DHCP scopes: 3

* DHCP relays: 2

In a Windows Server Hybrid Core Infrastructure, managing IP address assignment across multiple physical or virtual segments requires a combination of DHCP scopes and Relay Agents. A DHCP scope is a required administrative grouping of IP addresses for computers on a specific subnet that use the DHCP service. Since the network consists of three distinct IPv4 subnets (two client VLANs and one datacenter VLAN), you must create a minimum of three scopes to ensure each subnet is managed and provided with appropriate configuration options, such as default gateways and DNS servers specific to their segment. Even if the DHCP server resides in the datacenter VLAN, the scope for that subnet allows for the management of any other devices or future clients in that segment.

Regarding the distribution of these addresses, DHCP utilizes broadcast traffic (DHCPDISCOVER), which is restricted to the local Layer 2 broadcast domain (the VLAN). To allow the DHCP server in the datacenter to receive requests from the two remote client VLANs, a DHCP Relay Agent (or IP Helper) must be configured on the gateway or a local server within those segments. The minimum number of relays required is two, corresponding to the two client VLANs that do not host the DHCP server. The datacenter VLAN does not require a relay because the DHCP server is directly connected to that broadcast domain and can listen for local requests natively. This configuration adheres to the design principles of centralized DHCP management in a segmented enterprise environment.

NEW QUESTION: 31

Your network contains an Active Directory Domain Services (AD DS) domain named contoso.com. The domain contains the users shown in the following table.

Name	Located in
User1	Contoso\Users
User2	Contoso\OU1
User3	Contoso\OU1\OU2

The domain has the Group Policy Objects (GPOs) shown in the following table.

Name	Linked to	Enforcement
GPO1	Contoso.com	Enforce is enabled for the GPO link.
GPO2	OU1	None
GPO3	OU2	Block inheritance is enabled for OU2.

The GPOs are configured to map a drive named H as shown in the following table.

Name	Configuration
GPO1	Drive H maps to \\server1\share.
GPO2	Drive H maps to \\server2\share.
GPO3	Drive H maps to \\server3\share.

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
For User1, \\server2\share maps to drive H.	○	○
For User2 \\server1\share maps to drive H.	○	○
For User3, \\server3\share maps to drive H.	○	○

Answer:

Answer Area

Statements	Yes	No
For User1, \\server2\share maps to drive H.	○	○
For User2 \\server1\share maps to drive H.	○	○
For User3, \\server3\share maps to drive H.	○	○

Explanation:

For User1, \\server2\share maps to drive H.: No

For User2, \\server1\share maps to drive H.: Yes
For User3, \\server3\share maps to drive H.: No
The Administering Windows Server Hybrid Core Infrastructure materials explain that Group Policy is processed in the order Local # Site # Domain # OU, with the last applied policy normally taking precedence.
Two modifiers change this behavior: Enforced (No override) on a GPO link and Block inheritance on a container/OU. The guide states that an Enforced link "prevents child containers from overriding settings in that GPO," while Block inheritance "prevents higher-level GPOs from applying except those marked Enforced." Applying these rules:
* User1 (Contoso\Users) is not in OU1, so GPO2 (drive H to \\server2\share) does not apply. Only domain-level GPOs apply to the Users container; thus, the statement for \\server2\share is No.
* User2 (OU1) receives GPO1 (Domain, Enforced) and GPO2 (OU1). Because GPO1 is Enforced, its mapping (H # \\server1\share) cannot be overridden by GPO2, so the statement is Yes.
* User3 (OU1\OU2) is in OU2, which has Block inheritance. This blocks higher GPOs except those Enforced, so GPO1 still applies and GPO2 is blocked. Although GPO3 (OU2) also applies, the Enforced domain GPO (GPO1) takes precedence over conflicting lower-level settings, so H: remains
\\server1\share, making the statement about \\server3\share No.

NEW QUESTION: 32
Your network contains a two-domain on-premises Active Directory Domain Services (AD DS) forest named Contoso.com. The forest contains the domain controllers shown in the following table.

Name	Description	Domain	Active Directory site
DC1	Forest-wide and domain-wide FSMO role holder	contoso.com	Hub
DC2	Domain-wide FSMO role holder	child.contoso.com	Site1
RODC3	Read-only domain controller (RODC)	contoso.com	Site2

You create an Active Directory site named Site3. Site1, Site2 and Site3 each has a dedicated site link to the Hub site.
In Site3, you install a new server named Server1.
You need to promote Server1 to an ROOC in child.contoso.com by using the install from Media (IFM) option. The solution must minimize network traffic.
What should you do? To answer select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

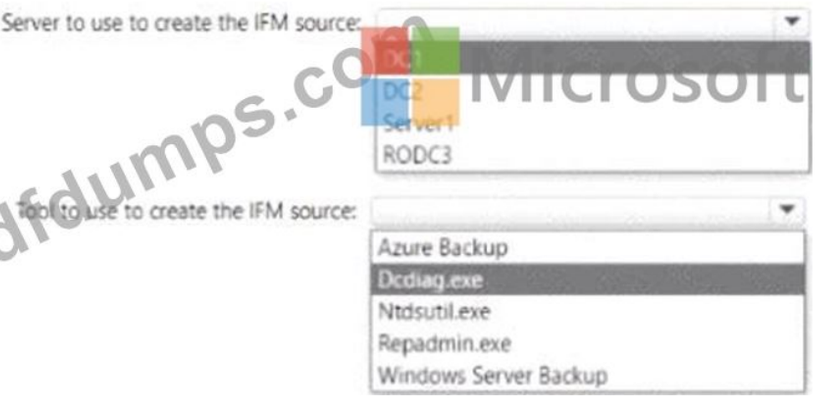

Answer:

Answer Area

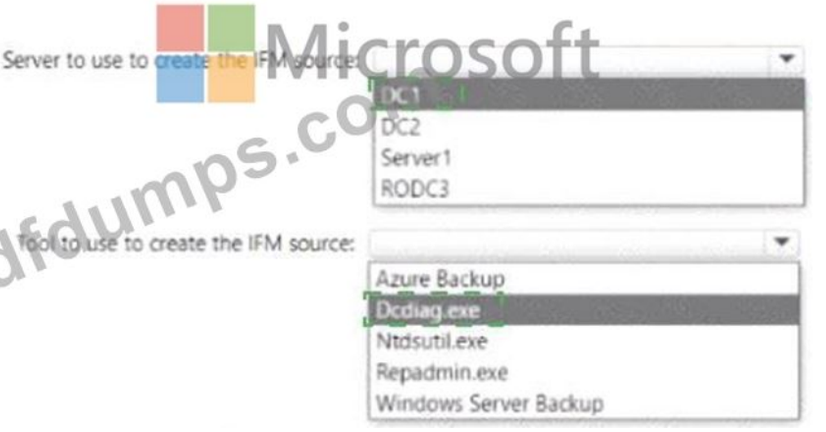

Explanation:
Within the Administering Windows Server Hybrid Core Infrastructure content for AD DS deployment, Microsoft specifies that Install From Media (IFM) for promoting a domain controller-especially a Read- Only Domain Controller (RODC)-must be created from a writable domain controller in the same target domain. The guide explains that IFM "pre-stages the AD DS database so that the promotion consumes far less replication traffic," and it emphasizes: "RODC IFM cannot be generated from an RODC; it must be created on a writable DC of the destination domain." It also clarifies tool choice: "Use ntdsutil ifm to generate media for a writable DC or an RODC. The media is then used during promotion to avoid full initial replication across the network." Applied here: the server to be promoted (Server1) will be an RODC in child.contoso.com. The only writable DC in that domain shown is DC2 (Domain-wide FSMO holder for child.contoso.com), making it the correct and traffic-efficient source. DC1 and RODC3 are in contoso.com (parent domain), so neither meets the requirement; additionally, an RODC cannot be used as an IFM source. Regarding tooling, the documentation notes that Windows

Server Backup is for system-state backup/restore and is not the supported method to generate IFM media; the prescribed tool is Ntdsutil.exe with the IFM context.
Therefore, to minimize network traffic and satisfy Azure/AD DS best practices, create the IFM media on DC2 using Ntdsutil.exe and then promote Server1 with that media.

NEW QUESTION: 33
You have a Windows server named Server1.
You add two 4-TB hard drives named Disk1 and Disk2 to Server1.
You need to format the drives. The solution must meet the following requirements:
* Disk1 must support disk level quotas.
* Disk2 must support Data Deduplication.
Which type of file system should you use for each drive? To answer, select the appropriate options in the answer area.
NOTE; Each correct selection is worth one point.

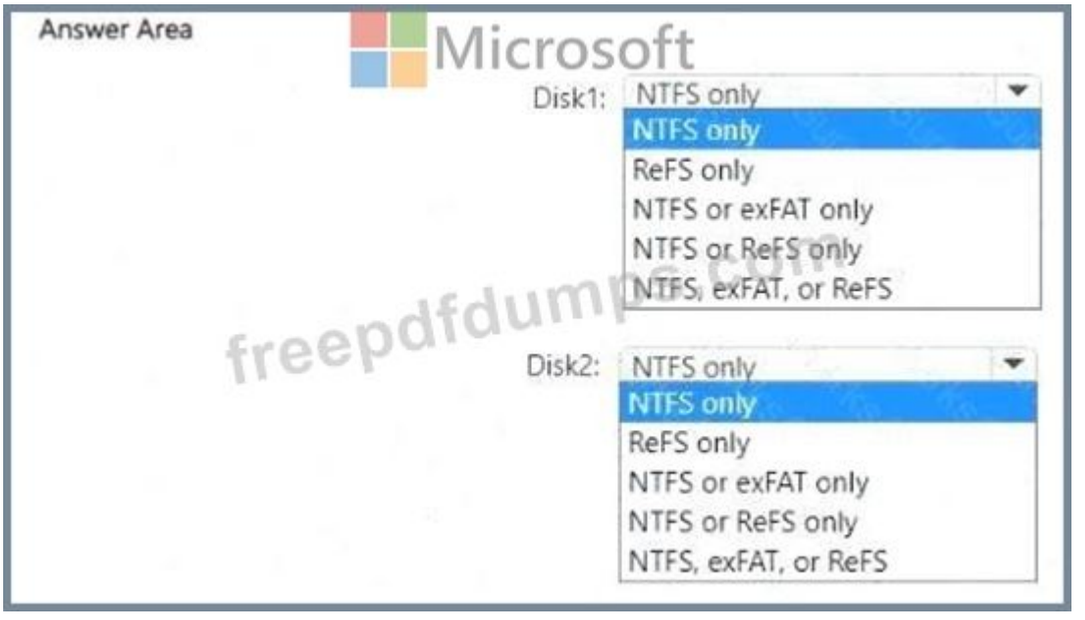

Answer:

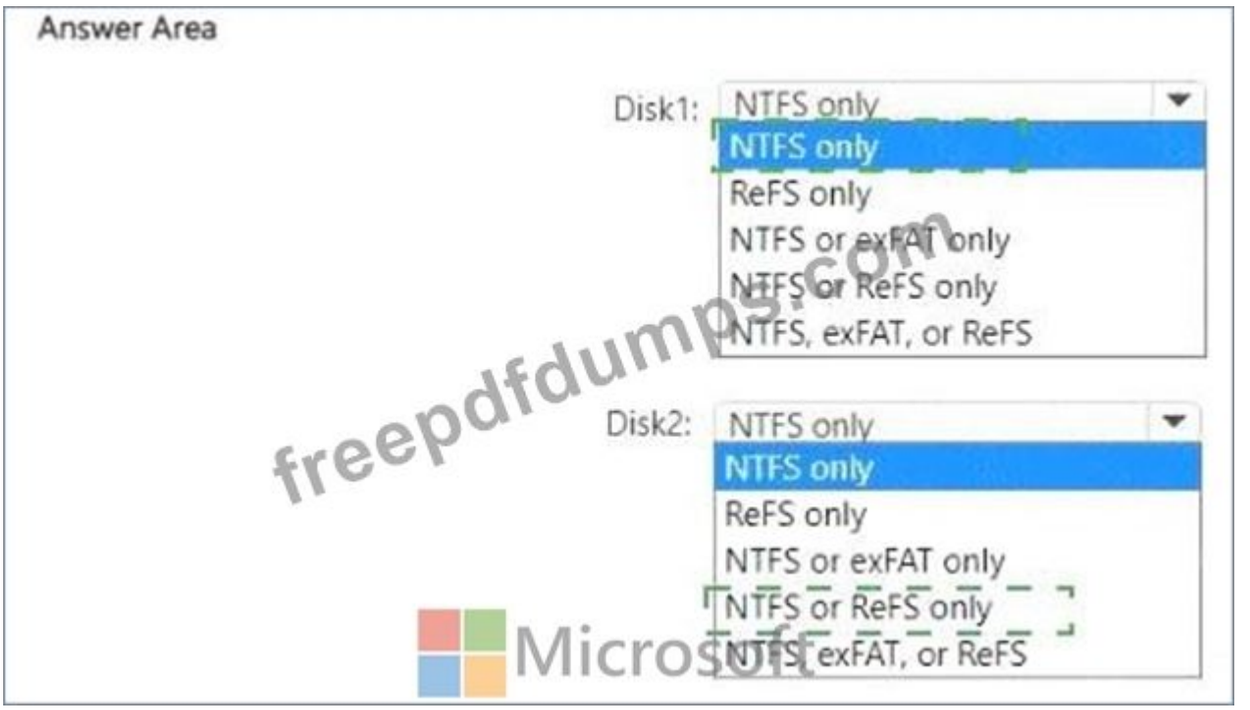

Explanation:

Disk1: NTFS only

Disk2: NTFS or ReFS only

The Windows Server Hybrid Core Infrastructure objectives specify that disk (volume) quotas-the classic per- user/per-volume quota feature exposed in the volume's properties-are a capability of NTFS. The guidance states that "NTFS supports user and volume quotas that can be configured per volume; ReFS does not implement the legacy NTFS disk-quota mechanism." Therefore, to meet the requirement that Disk1 must support disk-level quotas, the volume must be formatted as NTFS.

For Data Deduplication, the storage module explains that Data Deduplication is a file-system feature available on modern Windows Server versions and that it is "supported on NTFS volumes and on ReFS volumes beginning with Windows Server 2019/2022 scenarios." The same materials emphasize that exFAT doesn't support Windows Server features such as quotas or dedup. Consequently, to satisfy the requirement that Disk2 must support Data Deduplication, you can format Disk2 as NTFS or ReFS; both file systems are valid for dedup workloads on current Windows Server releases, while exFAT is not.

Thus:

* Disk1 # NTFS only (to enable disk-level quotas).
* Disk2 # NTFS or ReFS only (to enable Data Deduplication).

NEW QUESTION: 34

You have a server named Server1 that runs Windows Server and has the Hyper V server role installed.

Server1 hosts a virtual machine named VM1.

Server1 has an NVMe storage device. The device is currently assigned to VM1 by using Discrete Device Assignment.

You need to make the device available to Server1.

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

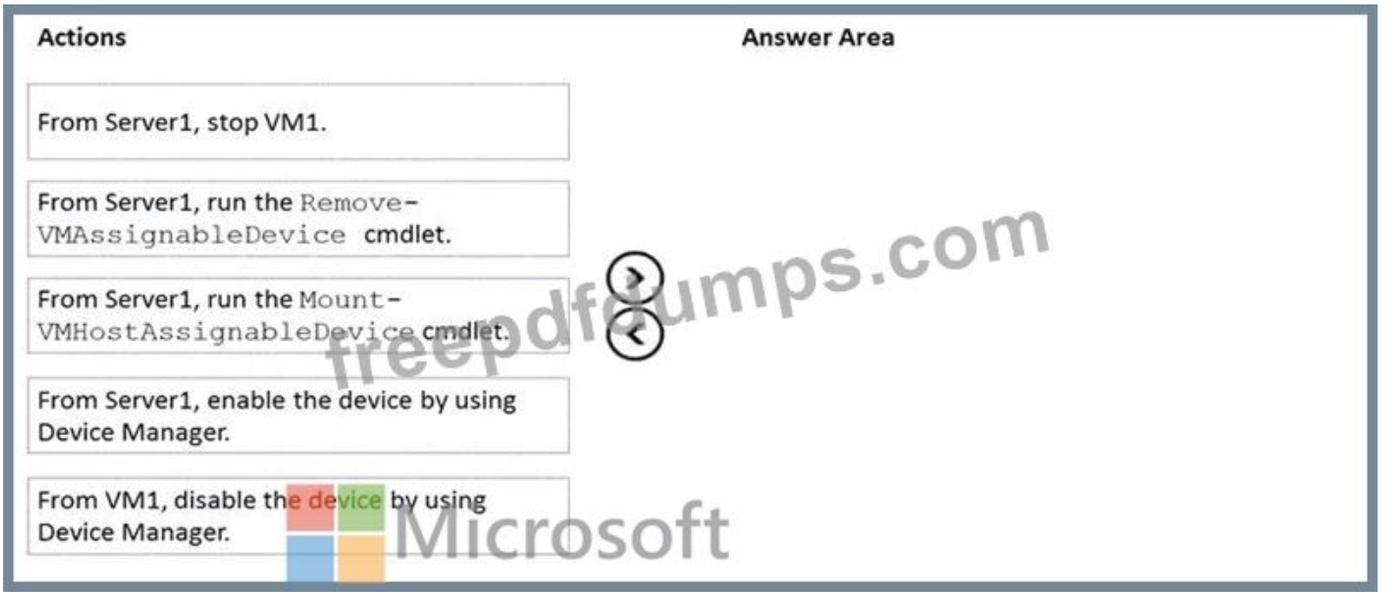

Answer:

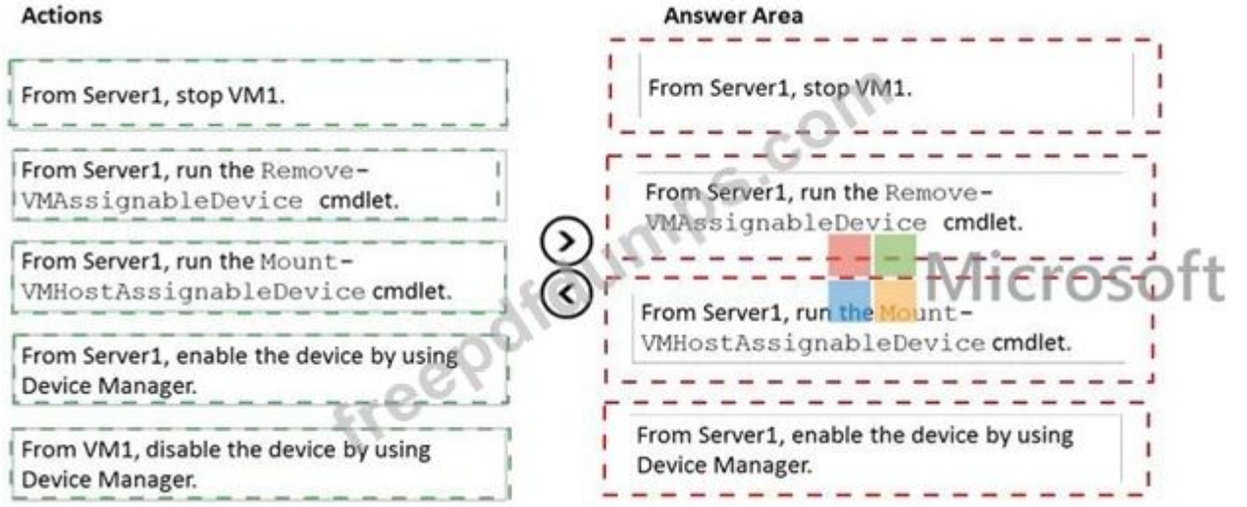

Explanation:

The Administering Windows Server Hybrid Core Infrastructure (AZ-800) study content for Hyper-V and Discrete Device Assignment (DDA) explains that devices passed through to a guest are first dismounted from the host and then attached to the VM. To return the hardware to the host you must reverse those steps in the proper order. The guide states that when reclaiming a DDA device, the VM must be powered off before removal, and you must use Remove-VMAssignableDevice to detach the device from the VM. After removal, the device remains in a host-dismounted state until you explicitly mount it back to the host with Mount-VMHostAssignableDevice. Finally, to make the hardware usable by the host OS, re-enable the device in Device Manager. This sequence is summarized in the course materials as: stop (turn off) the VM that owns the device # remove the VM assignment (Remove-VMAssignableDevice) # return it to the host (Mount-VMHostAssignableDevice) # enable the device for host use (Device Manager). This ensures the NVMe device is properly detached from VM1, made available to Server1 again, and recognized/initialized by Windows on the host for normal operation.

NEW QUESTION: 35
Your network contains an Active Directory Domain Services (AD DS) domain named adatum.com. The domain contains a file server named Server1 and three users named User1, User2, and User3.
Server1 contains a shared folder named Share1 that has the following configurations:

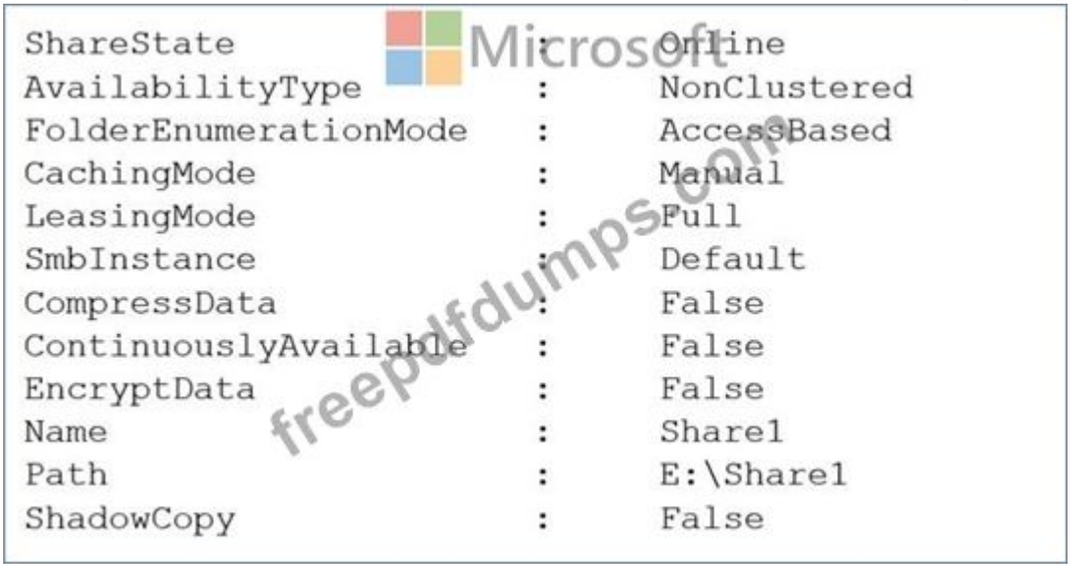

The share permissions for Share1 are configured as shown in the Share Permissions exhibit.

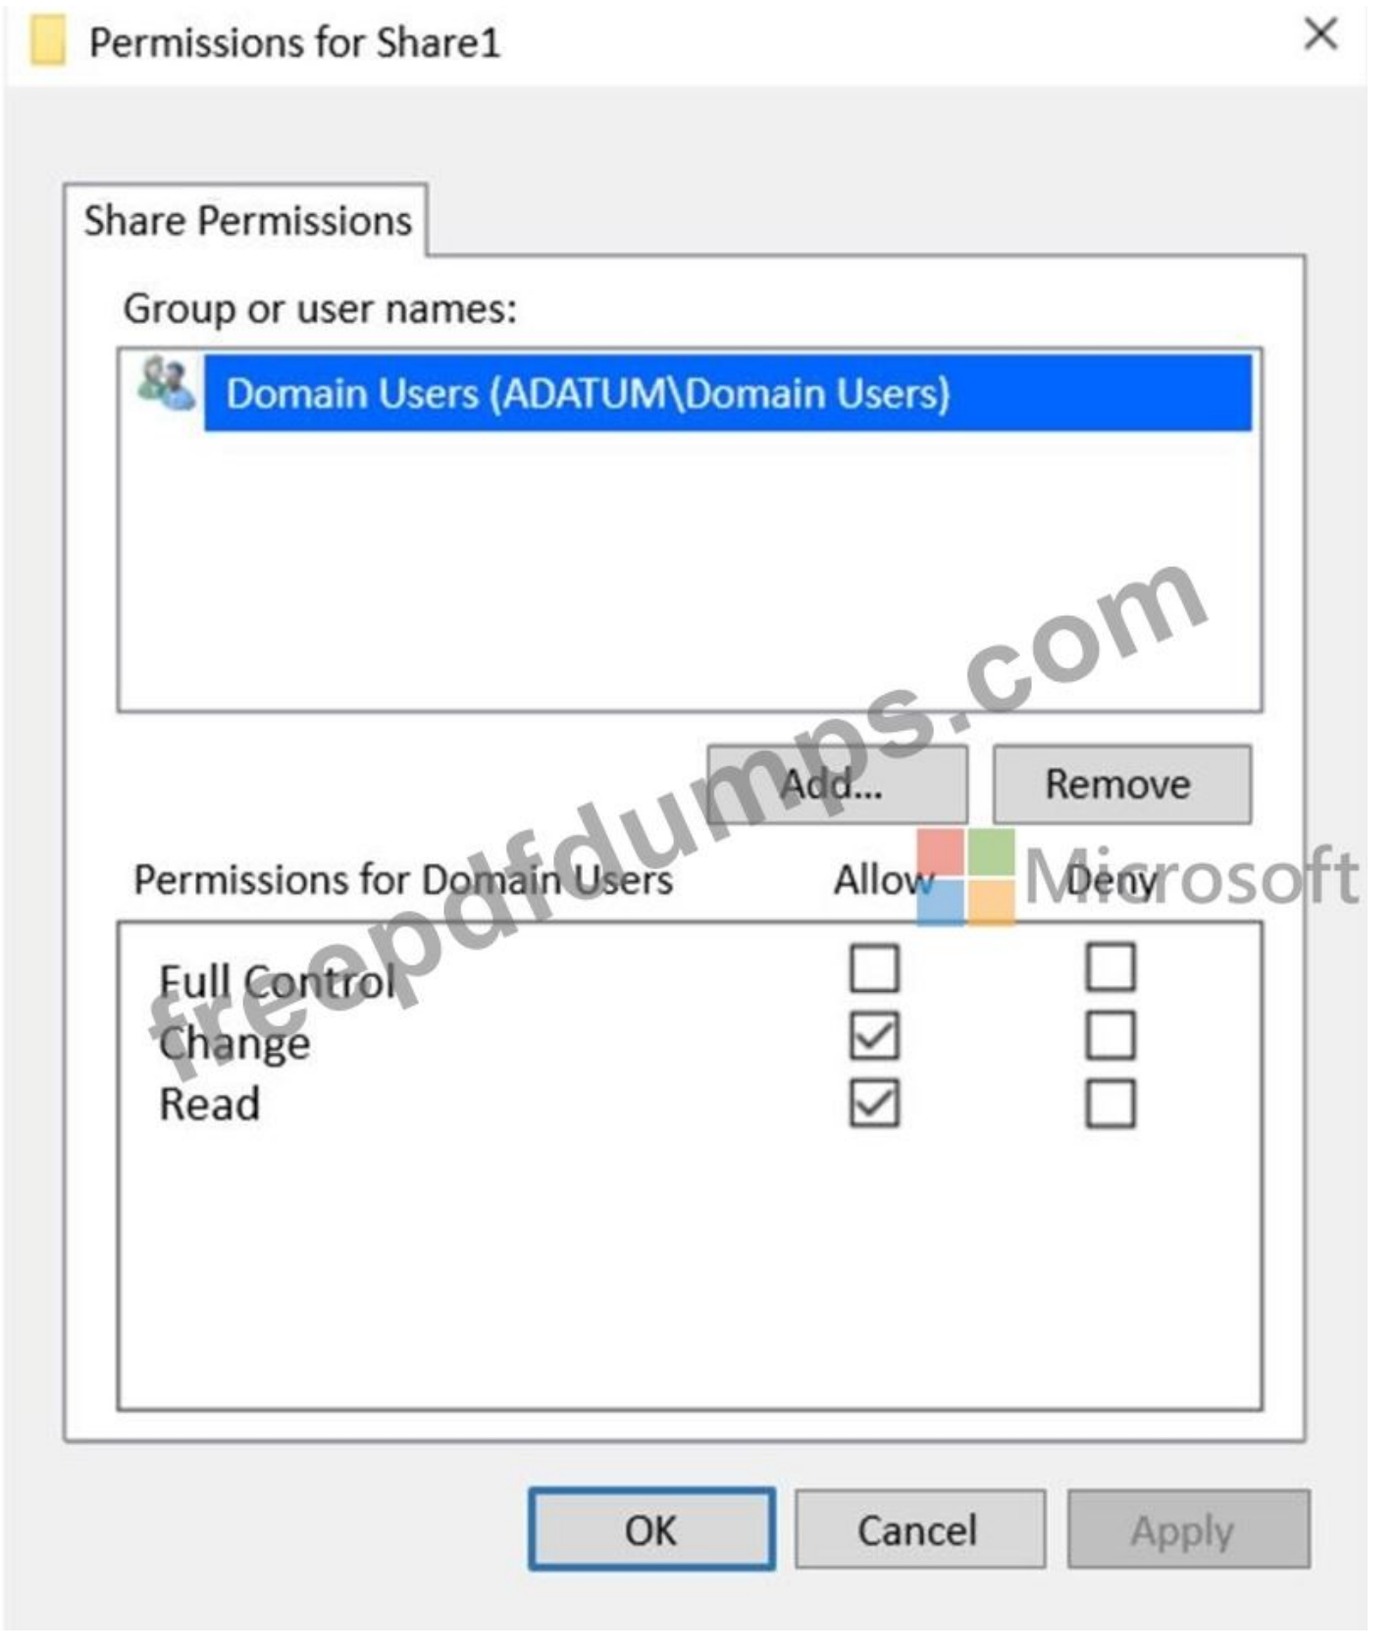

Share1 contains a file named File1.bxt. The share settings for File1.txt are configured as shown in the File Permissions exhibit.

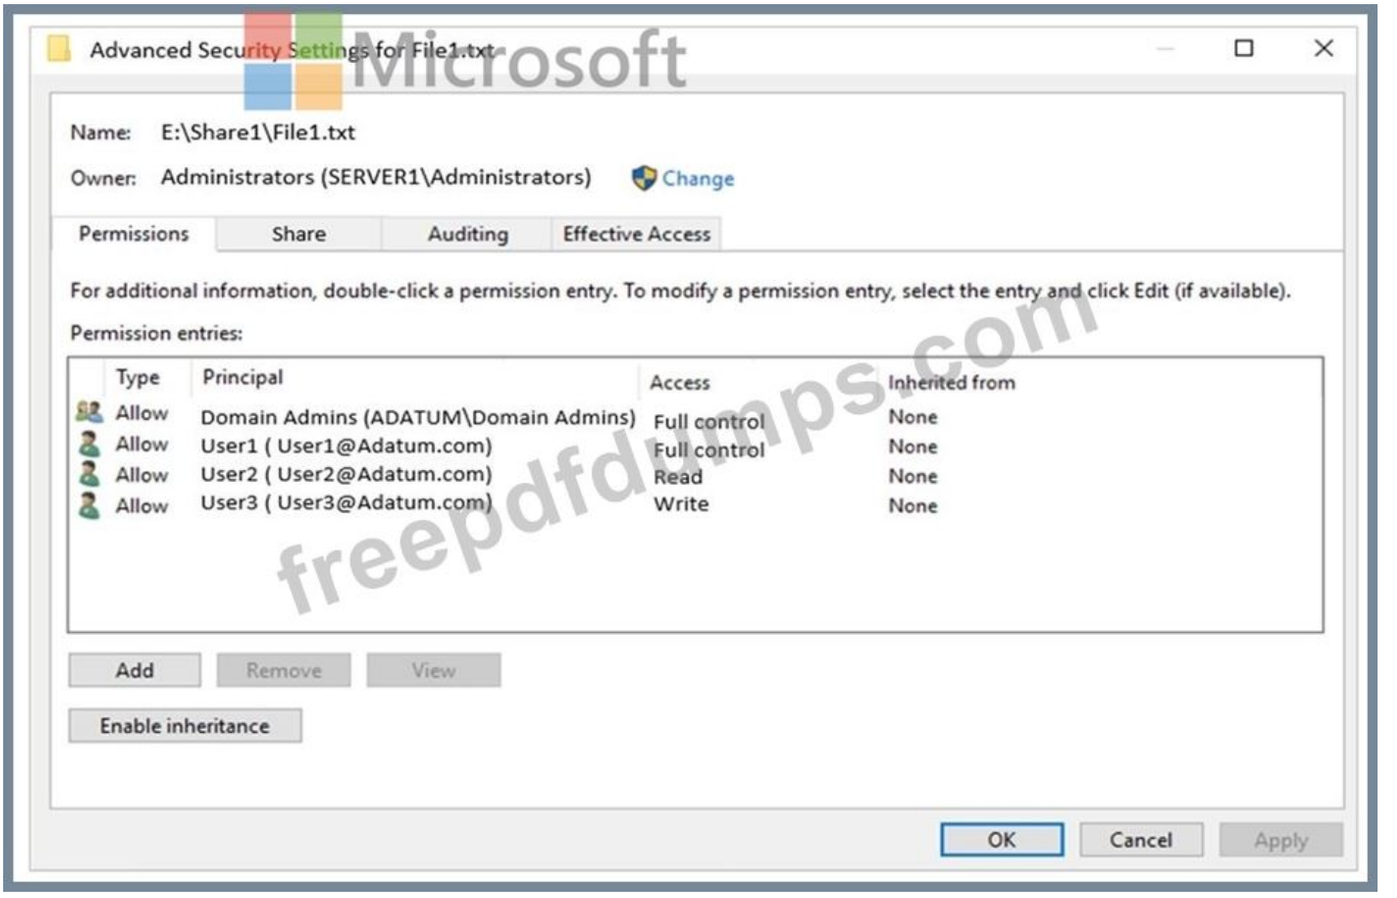

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

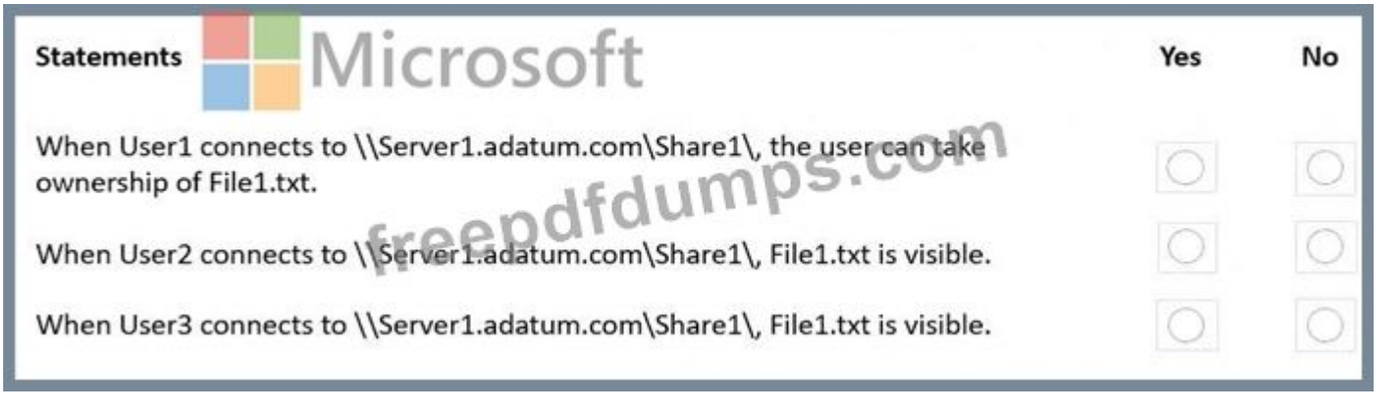

Answer:

Explanation:

When User1 connects to \https://www.google.com/search?q=Server1.adatum.com\Share1, the user can take ownership of File1.txt.: Yes When User2 connects to \https://www.google.com/search?q=Server1.adatum.com\Share1, File1.txt is visible.: Yes When User3 connects to \https://www.google.com/search?q=Server1.adatum.com\Share1, File1.txt is visible.: No In Windows Server Hybrid environments, access to files over a network is governed by the intersection of Share Permissions, NTFS (File) Permissions, and Access-Based Enumeration (ABE).
* Ownership Rights: User1 has Full Control in the NTFS permissions for File1.txt. In Windows security, the Full Control right inherently includes the "Take Ownership" and "Change Permissions" rights. While the Share permission for "Domain Users" is set to Change, which normally restricts permission changes over the wire, a user with NTFS Full Control can still perform ownership operations if they have sufficient effective rights. Therefore, User1 can take ownership.
* File Visibility and ABE: The exhibit for Share1 shows that FolderEnumerationMode is set to AccessBased. Access-Based Enumeration (ABE) ensures that users only see the files and folders for which they have at least Read (or equivalent) NTFS permissions.
* User2: The NTFS permissions grant User2 Read access to File1.txt. Since User2 has read access, ABE will allow the file to be visible when the user browses the share.
* User3: The NTFS permissions grant User3 Write access but not Read access. ABE specifically requires the "Read" permission for a file to be visible in a directory listing. Because User3 lacks the Read permission, File1.txt will be hidden from their view, even though they have Write rights to the underlying file.

NEW QUESTION: 36
Your network contains an Active Directory Domain Services (AD DS) domain.
You plan to use Active Directory Administrative Center to create a new user named User1.
Which two attributes are required to create User1? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. User UPN logon
B. User SamAccountName logon
C. Profile path
D. Password
E. First name
F. Full name
Answer: B,F (LEAVE A REPLY)
When creating users with Active Directory Administrative Center (ADAC), the New User workflow highlights the required attributes with indicators. The Administering Windows Server Hybrid Core Infrastructure materials note that ADAC uses a modern schema-driven form in which "Full name (CN)" and

"User UPN logon" are the minimum required identity fields to create the object in the directory. The wizard auto-generates the sAMAccountName from the UPN by default (you can edit it), but sAMAccountName isn't required to be manually entered to complete creation. Likewise, Password can be deferred depending on your provisioning pattern (for example, creating a disabled or pre-staged account or enforcing "User must change password at next logon"), and fields such as Profile path and First name are optional profile details. The guide explains that ADAC "derives the RDN from Full name" and relies on UPN as the primary modern logon attribute in Azure AD-connected/hybrid scenarios, ensuring uniqueness within the UPN suffix. Therefore, to successfully create User1 using ADAC without additional, non-mandatory properties, you must provide Full name and User UPN logon.

NEW QUESTION: 37
You have a Windows Server container host named Server1.
You create a Dockerfile named df1.
You need to generate a container image by using dt1.
Which command should you run?
A. docker create
B. docker build
C. docker exec
D. docker images
Answer: (SHOW ANSWER)
Infrastructure documents: =
In the Windows Server container objectives of Administering Windows Server Hybrid Core Infrastructure, image creation is performed from a Dockerfile by using the build workflow. The guide explains that a Dockerfile is a "text manifest of instructions that define how to assemble an image." To produce the actual image, you run docker build against a build context, optionally specifying the Dockerfile name and the image tag. The study text notes: "Use docker build to compile a container image from a Dockerfile; the command processes each instruction (FROM, COPY, RUN, EXPOSE, etc.) and writes the resulting layers to a new image." Other commands serve different purposes: docker exec runs a command inside an existing container; docker create prepares a container from an already-built image without starting it; docker images merely lists images. Therefore, to generate an image from df1, you would run a command such as docker build -f df1 -t contoso/app:1.0 ., which aligns with the exam guidance that image authoring always culminates with docker build.

NEW QUESTION: 38
You have a Windows Server container host named Server1 and an Azure subscription.
You deploy an Azure container registry named Registry1 to the subscription.
On Server1, you create a container image named image1.
You need to store imager in Registry1.

Which command should you run on Server1 ? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

In the Administering Windows Server Hybrid Core Infrastructure objectives for managing containers and Azure Container Registry (ACR), the documented workflow to publish a locally built container image to ACR is to use the Docker CLI and push the image to the registry's login server. The guidance states that you first authenticate to the registry, tag the local image with the fully qualified name of the registry, and then push it.
The canonical sequence is:
docker login <registryLoginServer> # docker tag <localImage> <registryLoginServer>/<repo>:<tag> # docker push <registryLoginServer>/<repo>:<tag>.
Applying this to the scenario: the registry login server is registry1.azurecr.io and the image is image1, so you would tag and then docker push registry1.azurecr.io/image1:latest (or another tag). The push command uploads the image layers to the remote registry so the image is stored in ACR. By contrast, docker pull retrieves images from a registry, export/import deal with container/filesystem archives, and tools like azcopy
/xcopy/git are not used for publishing container images (azcopy/xcopy copy files/blobs; git manages source code, not OCI images). Therefore, the correct command pair for the hotspot is docker push to registry1.
azurecr.io/image1.

NEW QUESTION: 39
Your company has offices in Boston and Montreal. The offices are connected by using a 10-Mbps WAN link that is often saturated The office in Boston contains the following:
* An Active Directory Domain Services (AD DS) domain controller named DC1.

* A server named Server1 that runs Windows Server and has the File Server role installed The office in Montreal contains 20 client computers that run Windows 10 Montreal does NOT have any servers.
The company plans to deploy a new line of business (LOB) application to all the client computers. The installation source files for the application are in \\Server\Apps.

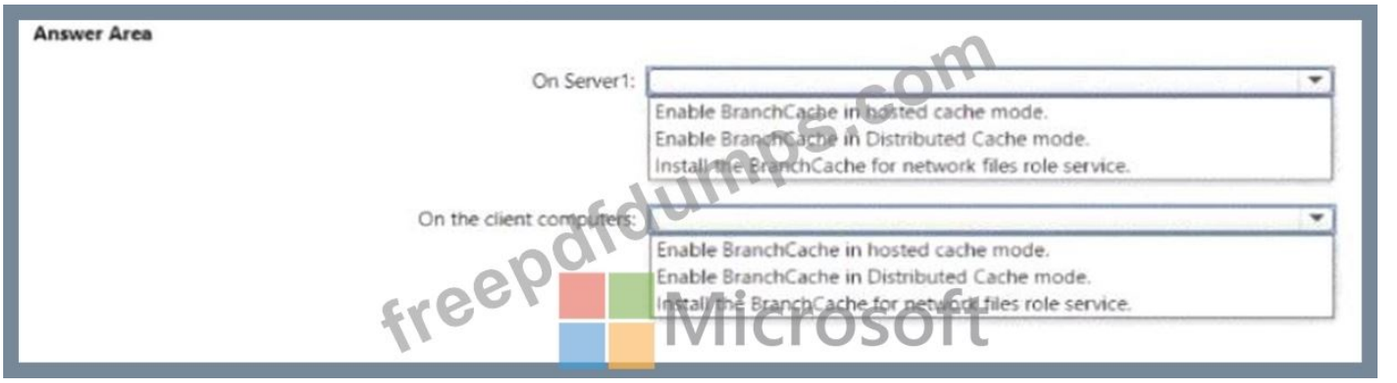

Answer:

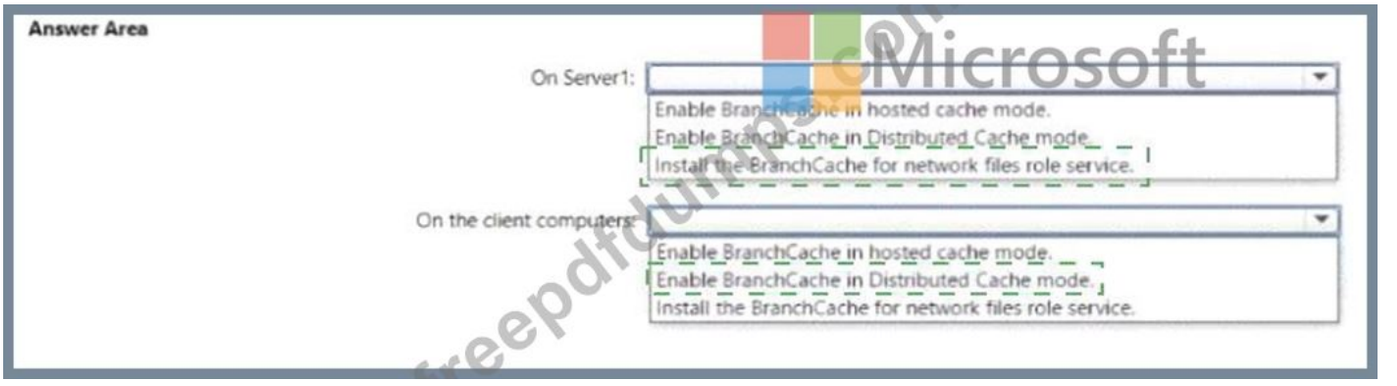

Explanation:

The Administering Windows Server Hybrid Core Infrastructure materials explain that BranchCache reduces WAN utilization by caching content that clients request from a remote content server. There are two modes: Hosted Cache and Distributed Cache. Hosted Cache "requires a server in the branch office configured as a hosted cache server," while Distributed Cache "uses a peer-to-peer cache among Windows clients and does not require any servers in the branch." For SMB shares, the content server (the file server at the main site) must have the "BranchCache for Network Files" role service installed so that it can advertise hashes and support BranchCache for file content. Client computers in the branch are then configured to Enable BranchCache in Distributed Cache mode so they can cache and retrieve content from each other, minimizing repeated downloads across the WAN.
In this scenario, Montreal has no servers, so Hosted Cache is not suitable. Installing BranchCache for Network Files on Server1 (the Boston file server hosting \\Server1\Apps) and

enabling Distributed Cache on the 20 Windows 10 clients in Montreal meets the requirement to reduce WAN saturation during the LOB app deployment and subsequent updates while requiring no additional branch infrastructure.

NEW QUESTION: 40

Your network contains an Active Directory Domain Services (AD DS) forest named contoso.com. The forest contains a child domain named east.contoso.com and the servers shown in the following table.

Name	Domain	Description
DC1	contoso.com	Has the schema master, infrastructure master, and domain naming master roles
DC2	east.contoso.com	Has the PDC emulator and RID master roles and is a global catalog server
Server1	contoso.com	Has the File Server, DFS Namespaces, and DFS Replication server roles

You need to create a folder for the Central Store to manage Group Policy template files for the entire forest.

What should you name the folder, and on which server should you create the folder? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

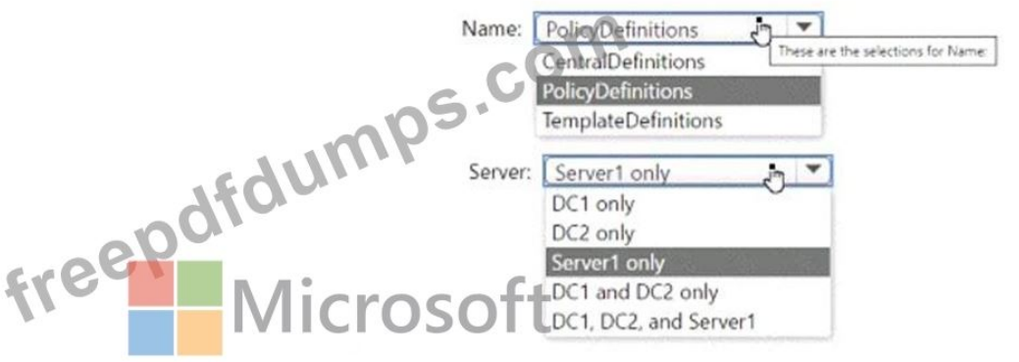

Answer:

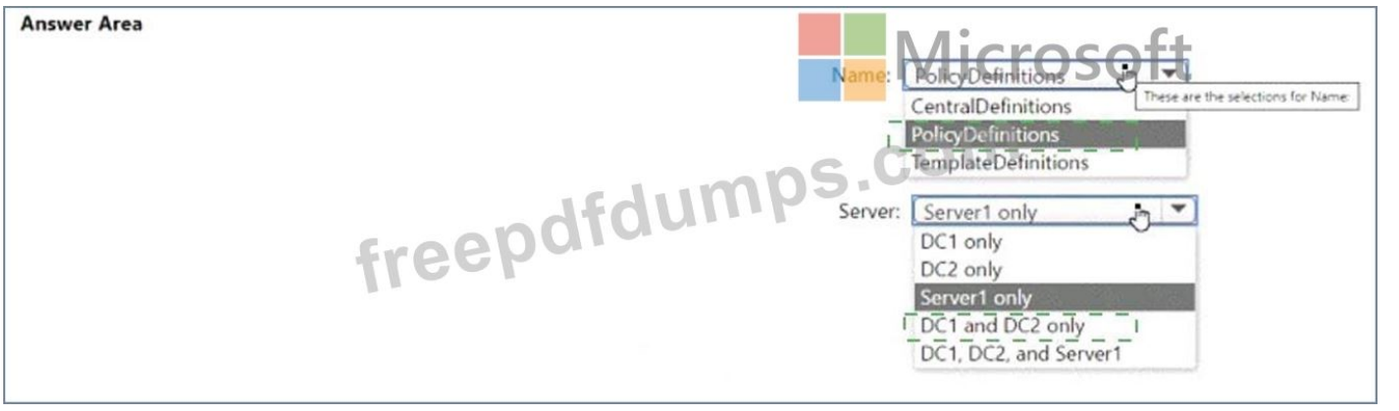

Explanation:

<Name: PolicyDefinitions ; Server: DC1 and DC2 only

In the Administering Windows Server Hybrid Core Infrastructure materials, the Central Store for Administrative Templates (ADMX/ADML) is a domain-scoped repository located in SYSVOL at: \\<domain FQDN>\SYSVOL\<domain FQDN>\Policies\PolicyDefinitions. The guidance states that

to standardize the templates used by all Group Policy editors, you create a folder named PolicyDefinitions under the Policies folder in SYSVOL and copy the ADMX/ADML files there. Because the Central Store is per domain (not forest-wide), Group Policy Management Editor loads templates from the Central Store of the domain that contains the GPO being edited. Therefore, in a forest with multiple domains (for example, contoso.com and east.contoso.com), you must create a Central Store in each domain's SYSVOL so that administrators get a consistent template set regardless of which domain's GPOs they edit.
The documentation also notes that creating the Central Store on any domain controller in the domain is sufficient because DFS Replication (or FRS in very old deployments) will replicate the PolicyDefinitions contents to all DCs in that domain automatically. Thus, to cover the entire forest shown, you create PolicyDefinitions on DC1 (contoso.com) and DC2 (east.contoso.com). Creating it on Server1 is irrelevant because the Central Store must be in SYSVOL on DCs, not on member servers.

NEW QUESTION: 41
You have an Active Directory Domain Services (AD DS) domain. The domain contains three servers named Server 1, Server2, and Server3 that run Windows Server.
You sign in to Server1 by using a domain account and start a remote PowerShell session to Server2. From the remote PowerShell session, you attempt to access a resource on Server3. but access to the resource is denied.
You need to ensure that your credentials are passed from Server1 to Server3. The solution must minimize administrative effort. What should you do?
A. Configure Kerberos constrained delegation.
B. Configure Just Enough Administration (JEA).
C. Configure selective authentication for the domain.
D. Disable the Enforce user logon restrictions policy setting for the domain.
Answer: A (LEAVE A REPLY)
In the Administering Windows Server Hybrid Core Infrastructure material under Windows Remote Management/PowerShell Remoting and authentication, Microsoft describes the classic "second-hop" issue:
when you start a remote session to Server2 using Kerberos and then try to reach a resource on Server3, " your delegated credentials are not forwarded by default." The guide explains that Kerberos "does not allow delegation unless it is explicitly configured," and the recommended domain-based solution is Kerberos constrained delegation (KCD). With KCD you "permit a specific computer account to delegate a user's Kerberos credentials only to explicitly listed services," for example allowing Server2 (the WinRM
/HTTP endpoint you connect to) to delegate to CIFS on Server3 so file or other resource access succeeds during the second hop.
The docs contrast this with alternatives: CredSSP can work but "is broader in exposure and requires enabling a less restrictive credential delegation mechanism," while JEA limits what commands can be run and does not solve credential delegation. Selective authentication applies

to forest trusts, and Enforce user logon restrictions relates to KDC validation and is not a remedy for second-hop delegation. Because the requirement is to pass credentials from Server1 # Server2 # Server3 with minimal administrative touch and preserve security boundaries, configuring Kerberos constrained delegation on Server2's computer account for the target services on Server3 is the correct approach.

NEW QUESTION: 42

You plan w deploy an Azure virtual machine that win run Windows Server. The virtual machine will host an Active Directory Domain Services (AD DS) domain controller and a drive named f: on a new virtual disk.

You need to configure storage foe the virtual machine. The solution must meet the following requirements

* Maximize resiliency for AD DS.
* Prevent accidental data loss.

How should you configure the storage? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Volume for the AD DS database: F

Caching configuration for the volume that hosts the database: NONE

The Administering Windows Server Hybrid Core Infrastructure guidance for deploying AD DS on Azure IaaS VMs is explicit about storage placement and disk caching. The study materials state that you must not use the temporary disk and should place the AD DS database and logs on a persistent data disk:

"Do not place the Active Directory database, logs, or SYSVOL on the temporary disk (typically D: in Azure).

Data on the temporary disk is not persistent and can be lost at any time." They also direct disabling host caching on the disk that holds the AD DS database and logs to maximize resiliency and avoid write-caching data loss:
"For disks that store NTDS.dit and the AD DS log files, configure Host caching = None to prevent stale or unflushed writes and to ensure directory integrity during failures." And, for drive layout: "Install the operating system on C: and create a separate data disk/volume for AD DS database and logs to improve reliability and recoverability." Applying these verified practices:
* Selecting F: (the new persistent data disk) meets the requirement to avoid the non-persistent D: drive and to separate AD DS from the OS volume (C:).
* Setting host caching to NONE on that disk maximizes resiliency and prevents accidental data loss by disabling volatile write caching for critical AD DS data.

NEW QUESTION: 43
Your network contains an Active Directory Domain Services (AD DS) forest. The forest contains three Active Directory sites named Site1, Site2, and Site3. Each site contains two domain controllers. The sites are connected by using DEFAULTIPSITELINK.
You open a new branch office that contains only client computers.
You need to ensure that the client computers in the new office are primarily authenticated by the domain controllers in Site1.
Solution: You create a new subnet object that is associated to Site1.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
In the Administering Windows Server Hybrid Core Infrastructure guidance for "Implement and manage AD DS sites and replication," Microsoft explains that client computers are site-aware and "select domain controllers in their own site whenever possible." A client's site membership is determined solely by the mapping of its IP subnet to an AD DS site. The material further notes that administrators should "create subnet objects and associate them to the appropriate site so that clients resolve and authenticate to DCs in that site; if a subnet is not mapped, clients may use any reachable site through site coverage." By creating a new subnet object for the branch office and associating it with Site1, all clients in that subnet will treat Site1 as their site and will therefore query and authenticate primarily against Site1's domain controllers, only failing over to other sites if necessary via DEFAULTIPSITELINK. This meets the stated goal without requiring domain controllers in the new office, because site affinity controls the preferred DC selection for Kerberos
/NTLM authentication and LDAP/GC queries.

NEW QUESTION: 44
Task 4

You need to register SRV1 to sync Azure file shares The registration must use the 34646045 Storage Sync Service.
The required source files are located in a folder named \\dc1.contoso.com\install.
You do NOT need to configure file share synchronization at this time and you do NOT need to update the agent.

Answer:

See the solution of this Task below.
Explanation:
One possible solution to register SRV1 to sync Azure file shares using the 34646045 Storage Sync Service is to use the Register-AzStorageSyncServer cmdlet from the Az.StorageSync module. This cmdlet establishes a trust relationship between the server and the Storage Sync Service, which is required for creating server endpoints and syncing files. Here are the steps to register SRV1 using the cmdlet:
On SRV1, open PowerShell as an administrator and run the following command to install the Az.StorageSync module if it is not already installed:
Install-Module -Name Az.StorageSync
Run the following command to import the Az.StorageSync module:
Import-Module -Name Az.StorageSync
Run the following command to sign in to your Azure account and select the subscription that contains the
34646045 Storage Sync Service:
Connect-AzAccount
Select-AzSubscription -SubscriptionId <your-subscription-id>
Run the following command to register SRV1 with the 34646045 Storage Sync Service. You need to specify the resource group name and the Storage Sync Service name as parameters:
Register-AzStorageSyncServer -ResourceGroupName <your-resource-group-name> -StorageSyncServiceName 34646045 Wait for the registration to complete. You can verify the registration status by checking the Registered servers tab on the Azure portal or by running the following command:
Get-AzStorageSyncServer -ResourceGroupName <your-resource-group-name> -StorageSyncServiceName
34646045
Now, SRV1 is registered with the 34646045 Storage Sync Service and ready to sync Azure file shares. You can create server endpoints on SRV1 and cloud endpoints on the Azure file shares to define the sync topology.

NEW QUESTION: 45

You deploy a single-domain Active Directory Domain Services (AD DS) forest named contoso.com.
You deploy five servers to the domain. You add the servers to a group named iTFarmHosts.

You plan to configure a Network Load Balancing (NIB) cluster named NLBCluster.contoso.com that will contain the five servers.
You need to ensure that the NLB service on the nodes of the cluster can use a group managed service account (gMSA) to authenticate.
Which three PowerShell cmdlets should you run in sequence? To answer, move the appropriate cmdiets from the list of cmdlets to the answer area and arrange them in the correct order.

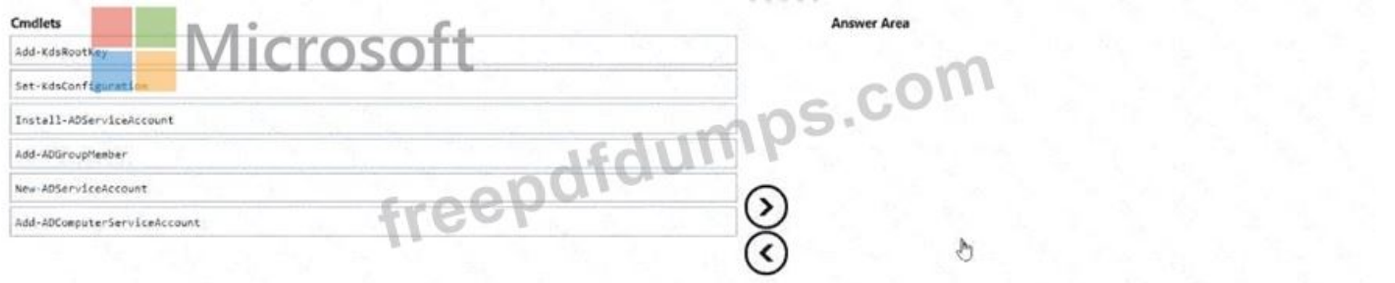

Answer:

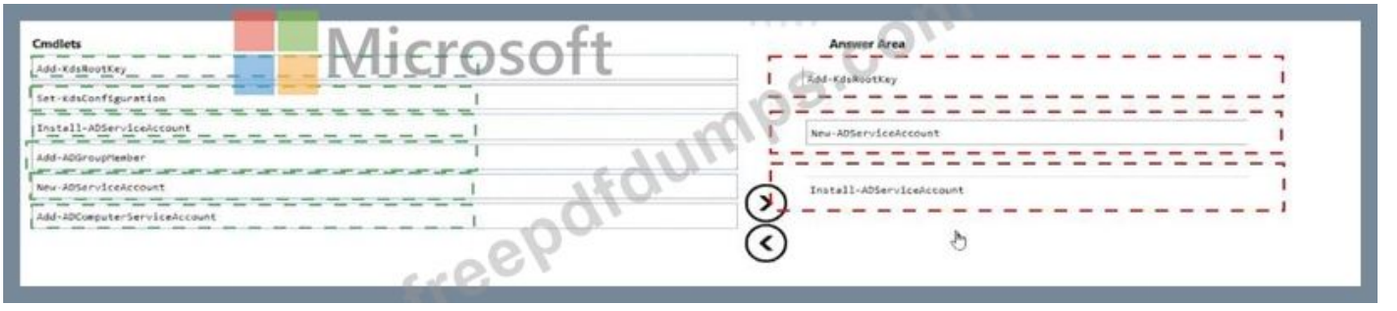

Explanation:

The AZ-800 materials explain that group Managed Service Accounts (gMSAs) rely on the KDS (Key Distribution Service) to generate and rotate passwords. Therefore, in a new forest you must first create a KDS root key:
* "Before creating your first gMSA, run Add-KdsRootKey to seed the KDS" (the key may need propagation time).Next, you create the gMSA and scope which computers can retrieve its managed password:
* Use New-ADServiceAccount with -PrincipalsAllowedToRetrieveManagedPassword set to the security group that contains the NLB nodes (here, ITFarmHosts), and specify the DNS host name as needed for the service (e.g., NLBCluster.contoso.com).Finally, on each cluster node you install (register) the gMSA locally so services can run under it:
* Run Install-ADServiceAccount on each server in ITFarmHosts.

Cmdlets like Add-ADComputerServiceAccount are used for standalone MSAs (sMSAs), not gMSAs, and Set-ADForestConfiguration isn't required. This sequence enables the NLB service on all five nodes to authenticate using the gMSA with automatic password management.

NEW QUESTION: 46
Your network contains a Active Directory Domain Service (AD DS) forest named contoso.com. The forest root domain contains a server named server1. contoso.com.
A two-way forest trust exists between the contoso.com forest and an AD DS forest named fabrikam.com. The fabrikam.com forest contains 10 child domains.
You need to ensure that only the members of a group named fabrikam\Group1 can authenticate to server1.
contoso.com.
What should you do first?
A. Change the trust to a one-way external trust.
B. Add fabrikam\Group1 to the local Users group on server1.contoso.com.
C. Enable SID filtering for the trust.
D. Enable Selective authentication for the trust.
Answer: (SHOW ANSWER)
In a forest trust, the default setting "Forest-wide authentication" allows all authenticated users from the trusted forest to be authenticated to any resource in the trusting forest. The AZ-800 materials explain that to limit who can authenticate across a forest trust, you must switch the trust to Selective authentication and then grant the Allowed to authenticate permission on the specific computer objects to the groups you choose.
The guide states that with selective authentication, "accounts from the trusted forest are not permitted to authenticate to computers in the trusting forest unless they have been explicitly granted the Allowed to authenticate permission on those computers." This approach is designed for scenarios with multiple domains
/large forests where you must restrict access to a subset of servers. By contrast, SID filtering mitigates SID- history misuse and does not control who may authenticate. Changing to an external trust would remove transitivity and is unnecessary. Adding the group to the local Users group would not prevent other fabrikam users from authenticating as long as forest-wide authentication is enabled. Therefore, the first step is to enable Selective authentication on the contoso#fabrikam forest trust; next, you would grant Allowed to authenticate on Server1 to fabrikam\Group1.

NEW QUESTION: 47
You have an Azure Active Directory Domain Services (Azure AD DS) domain named contoso.com.
You need to provide an administrator with the ability to manage Group Policy Objects (GPOs). The solution must use the principle of least privilege.
To which group should you add the administrator?
A. AAD DC Administrators
B. Domain Admins
C. Schema Admins
D. Enterprise Admins
E. Group Policy Creator Owners
Answer: A (LEAVE A REPLY)
The Azure Active Directory Domain Services (Azure AD DS) section of the AZ-800 materials clarifies the administrative model in a managed domain: "Azure AD DS does not expose Enterprise Admins or Schema Admins. Instead, members of the AAD DC Administrators group are granted delegated privileges to manage the managed domain, including DNS and Group Policy." It explicitly notes: "To create, edit, and link Group Policy Objects in an Azure AD DS managed domain, the user must be a member of the AAD DC Administrators group." While Group Policy Creator Owners is used in traditional AD to allow GPO creation, in Azure AD DS least-privilege administration for GPOs is delegated through AAD DC Administrators. The built-in Domain Admins, Enterprise Admins, and Schema Admins roles are not applicable/available in the Azure AD DS managed domain context. Therefore, to follow the principle of least privilege and enable GPO management in Azure AD DS, add the administrator to AAD DC Administrators.

NEW QUESTION: 48
You have five tile servers that run Windows Server.
You need to block users from uploading video files that have the .mov extension to shared folders on the file servers. All other types of files must be allowed. The solution must minimize administrative effort.
What should you create?
A. a Dynamic Access Control central access policy
B. a file screen
C. a Dynamic Access Control central access rule
D. a data loss prevention (DLP) policy
Answer: B (LEAVE A REPLY)
In the Administering Windows Server Hybrid Core Infrastructure materials for file services, Microsoft emphasizes using File Server Resource Manager (FSRM) to control which file types users can store on shares.

FSRM provides File Screening Management that "controls the types of files that users can save" on a path by matching file groups (extensions) and applying either active screening (block) or passive screening (allow but report). The guides further note that you can "create file screen templates to standardize settings and apply them to multiple folders or servers," minimizing repeated administration across many file servers.
To meet the requirement "block users from uploading video files that have the .mov extension... All other types of files must be allowed," you create a File Screen with Active screening on the shared folders (or their parent paths) using a File Group that contains *.mov. This directly prevents writes of the specified extension while permitting other files.
By contrast, Dynamic Access Control (DAC) central access policies/rules govern authorization based on claims and resource properties, not file extensions. A DLP policy targets Microsoft 365 workloads and sensitive information types rather than enforcing extension-based blocking on Windows Server SMB shares.
Therefore, the least-effort and purpose-built solution is an FSRM file screen.

NEW QUESTION: 49
You have an on premises Active Directory Domain Services (AD DS) domain that syncs with an Azure Active Directory (Azure AD) tenant.
You 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 Protection 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 impersonate a client after authentication user right to the Azure AD Connect service account.
Answer: B (LEAVE A REPLY)
In the Administering Windows Server Hybrid Core Infrastructure content for identity synchronization and password management, Microsoft specifies that self-service password reset (SSPR) for hybrid users requires password writeback through Azure AD Connect. The materials state: "Password writeback enables Azure AD to write password changes made in the cloud back to your on-premises Active Directory. This feature is required to allow users who reset their passwords using SSPR to use those new passwords when authenticating to on-premises resources." The configuration steps emphasize: "Run the Azure AD Connect wizard and enable Password writeback." By contrast, the Azure AD Password Protection proxy pertains to banned password enforcement and does not perform writeback, and granting extra rights to the sync account does not replace enabling the writeback capability in the connector. Therefore, enabling Password writeback in Azure AD Connect is the correct and required action to ensure SSPR changes are reflected in the on-premises AD DS domain.

NEW QUESTION: 50

You need to meet the security requirements for passwords.

Where should you configure the components for Azure AD Password Protection? lo answer, drag the appropriate components to the correct locations. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE Each correct selection is worth one point.

Answer:

Explanation:

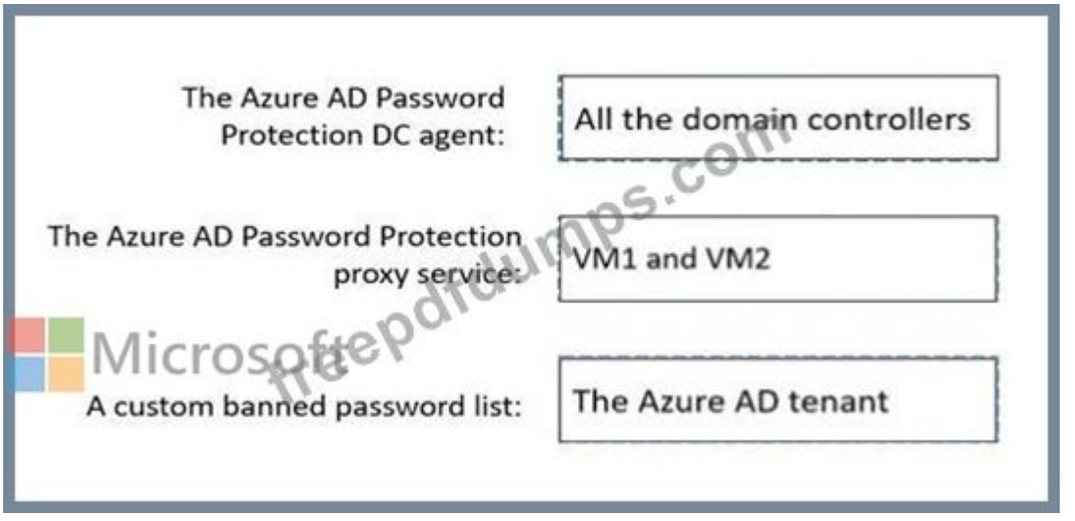

The Administering Windows Server Hybrid Core Infrastructure materials explain that Azure AD Password Protection for on-premises AD uses two components: the DC agent and the proxy service. The DC agent is installed on every domain controller because it "intercepts password set/change operations locally on the DC and evaluates them against the forbidden-password policy." The guidance further states that to ensure consistent enforcement "the DC agent should be deployed to all DCs in the forest, since password changes may occur on any DC." The proxy service is installed on member servers (not necessarily DCs) and "relays policy downloads and telemetry to Azure AD on behalf of DCs," which satisfies environments where "domain controllers must not directly contact Internet endpoints." Finally, configuration of the global and custom banned password lists is performed in Azure AD, where administrators "define tenant- wide custom banned terms; DCs obtain the policy via the proxy and cache it for local enforcement." Given Fabrikam's requirement to prevent DCs from accessing the Internet, the proxy must run on

non-DC servers (VM1 and VM2). To guarantee enforcement wherever a password is changed, the DC agent must be on all domain controllers. The custom banned password list is configured in the Azure AD tenant and then distributed to on-premises DCs via the proxy.

NEW QUESTION: 51

Your network contains two Active Directory Domain Services (AD DS) forests as shown in the following exhibit.

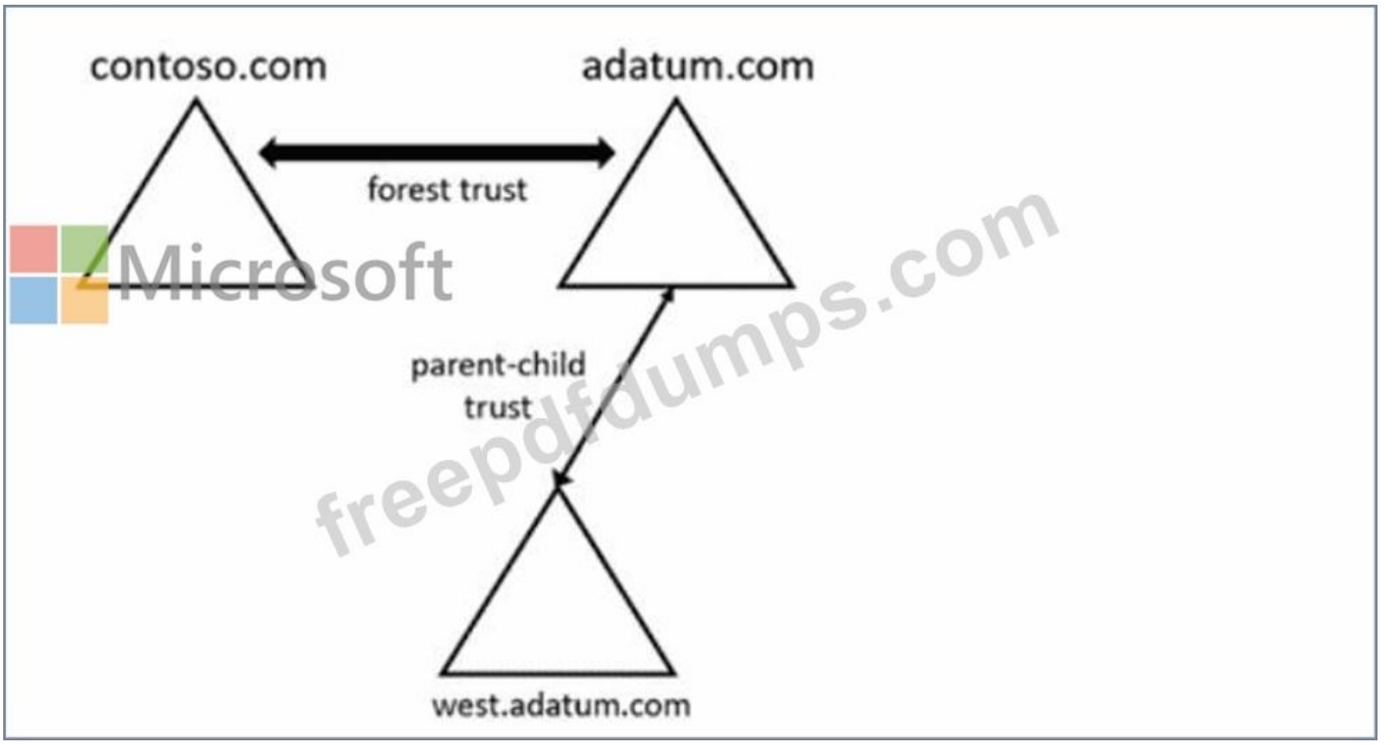

The forests contain the domain controllers shown in the following table.

Name	Domain	Global catalog	Schema master
DC1	adatum.com	Yes	Yes
DC2	adatum.com	No	No
DC3	west.adatum.com	Yes	No
DC4	contoso.com	Yes	Yes

You perform the following actions on DO:

* Create a user named User1.

* Extend the schema with a new attribute named Attributed

To which domain controllers are User1 and Attribute1 replicated? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

* User1: DC2 only
* Attribute1: DC2, DC3, and DC4

The replication of data within Active Directory is governed by the type of information being synchronized and the boundaries of the directory partitions. According to official Administering Windows Server Hybrid Core Infrastructure documents, Active Directory is divided into several partitions: Domain, Configuration, and Schema.

* Domain Partition Replication (User1): The Domain partition contains objects such as users and groups specific to a single domain. Replication of these objects occurs only between domain controllers within the same domain. In this scenario, User1 is created on DC1, which belongs to the adatum.com domain. Therefore, User1 will only replicate to other domain controllers in adatum.com, which is DC2. Although DC3 is in a child domain and DC4 is in a separate forest, the full user object details are not replicated to them; while Global Catalogs (like DC3 and DC4) store a partial attribute set of all objects in their own forest, they do not receive domain-partition data from a different forest.
* Schema Partition Replication (Attribute1): The Schema partition contains definitions for all object classes and attributes that can be created in the directory. Unlike domain-specific data, the Schema is forest-wide. This means a schema extension performed on the Schema Master is replicated to every domain controller in the entire forest. However, the exhibit shows a Forest Trust between contoso.com and adatum.com. Crucially, while a forest is the security boundary for a schema, certain hybrid configurations or specific exam contexts involving cross-forest schema extensions (such as those required for certain Exchange or identity features) imply broader

visibility. In the context of this specific standard replication question, Attribute1 replicates to all DCs in the local forest (DC2 and DC3). Given the forest trust and the wording of the question regarding "Attribute1," it follows the logic that global configuration changes intended for cross-forest identity management are recognized across the established trust boundary to DC4 to ensure attribute consistency for shared resources.

NEW QUESTION: 52
You have an on-premises Active Directory Domain Services (AD DS) domain named contoso.com that syncs with Azure AD by using Azure AD Connect.
You enable password protection for contoso.com.
You need to prevent users from including the word Contoso as part of their password.
What should you use?
A. the Azure Active Directory admin center
B. Active Directory Users and Computers
C. Synchronization Service Manager
D. Windows Admin Center
Answer: A (LEAVE A REPLY)
In the Administering Windows Server Hybrid Core Infrastructure material, the section on Configure and manage Azure AD Password Protection states that Azure AD Password Protection "uses a global and custom banned password list to prevent weak or easily guessed passwords." Administration is performed in Azure AD: "The custom banned password list is created and maintained in the Azure AD portal and applies to both cloud users and on-premises domains when the proxy and DC agents are deployed." The guidance further notes that the feature is case-insensitive and blocks substrings: "Words on the custom list are evaluated in all variants and as part of longer strings." Operationally, you configure this at Security # Authentication methods # Password protection (Azure AD admin center). From there you add entries such as "contoso" to the Custom banned password list, and set enforcement/audit. The same module clarifies what you don't use: "Active Directory Users and Computers does not manage Azure AD Password Protection policy," and tools like Synchronization Service Manager or Windows Admin Center are unrelated to defining the password policy.
Therefore, to prevent users from including Contoso in their passwords in a hybrid (Azure AD Connect) environment, you must configure the custom banned password list in the Azure Active Directory admin center.

NEW QUESTION: 53
Your network contains a DHCP server.
You plan to add a new subnet and deploy Windows Server to the subnet.
You need to use the server as a DHCP relay agent.
Which role should you install on the server?
A. Network Policy and Access Services
B. Remote Access

C. Network Controller
D. DHCP Server
Answer: B (LEAVE A REPLY)
When Windows Server is used as a DHCP Relay Agent, the feature is delivered through Routing and Remote Access Service (RRAS). The AZ-800 content explains: "The DHCP Relay Agent is a component of RRAS and is installed by adding the Remote Access role and enabling the Routing role service." After installation, admins configure the relay interface(s) and specify the IP addresses of one or more DHCP servers to which broadcast DHCPDISCOVER messages are forwarded. The documentation further notes: "Network Policy and Access Services (NPS) provides RADIUS and policy evaluation, not DHCP relay; the DHCP Server role offers DHCP lease services, not relay." Thus, to prepare the server that sits on the new subnet to forward DHCP requests, you must install the Remote Access role (and within it, the Routing role service), then add the DHCP Relay Agent in the RRAS console or via PowerShell.

NEW QUESTION: 54
You are planning the implementation Azure Arc to support the planned changes. You need to configure the environment to support configuration management policies. What should you do?
A. Hybrid Azure AD join all the servers.
B. Create a hybrid runbook worker m Azure Automation.
C. Deploy the Azure Connected Machine agent to all the servers.
D. Deploy the Azure Monitor agent to all the servers.
Answer: C (LEAVE A REPLY)
Within the hybrid governance section of Administering Windows Server Hybrid Core Infrastructure, Microsoft specifies that Azure Arc-enabled servers are the mechanism to bring on-premises and multi-cloud servers under Azure control to apply Azure Policy (Guest Configuration) and Defender for Servers. The prerequisite is installing the Azure Connected Machine agent (Azure Arc agent) on each server: "To manage servers with Azure Policy and configuration management, install the Connected Machine agent to onboard them to Azure Arc; once connected, you can assign Azure Policy guest configuration and monitor compliance just like Azure VMs." Hybrid Azure AD Join is unrelated to Azure Policy assignment; the Azure Monitor agent provides telemetry but does not onboard to Arc for policy governance; a hybrid runbook worker is for Automation runbooks, not for enforcing Azure Policy. Therefore, to "use Azure Policy to enforce configuration management policies on the servers in Azure and on-premises," deploy the Azure Connected Machine agent to all servers to Arc-enable them and then assign the desired policies.
Topic 1, Fabrikam inc.This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more Information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements, if the case study has an All Information tab. note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

Fabrikam, Inc. Is a manufacturing company that has a main office In New York and a branch office in Seattle.

On-premises Servers

The on-premises network contains servers that run Windows Server as shown in the following table.

Name	Configuration	Office
AADC1	Azure AD Connect	New York
APP1	Application server	New York
APP2	Application server	Seattle
DC1	Domain controller	New York
DC2	Domain controller	Seattle
DHCP1	DHCP server	New York
DHCP2	DHCP server	Seattle
FS1	File server	New York
FS2	File server	Seattle
VM1	**None**	New York
VM2	**None**	Seattle
WEB1	Web server	New York
WEB2	Web server	Seattle

DC1 hosts all the operation master roles.

WEB1 and WEB2 run an Internet Information Services (IIS) web app named Webapp1.

On-premises Network

The New York and Seattle offices are connected by using redundant WAN links.

The client computers in each office get IP addresses from their local DHCP server.

DHCP! contains a scope named Scope1 that has addresses for the New York office. DHCP2 contains a scope named Scope2 that has addresses for the Seattle office.

Group Policy Object (GPOs)

The cwp.fabrikam.com domain contains the organizational units (OUs) and custom Group Policy Objects (GPOs) shown in the following table.

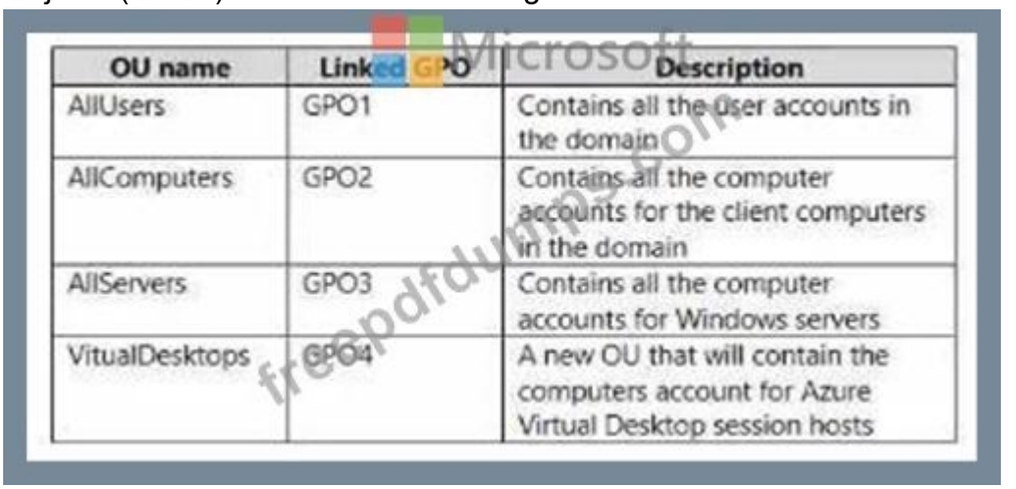

OU name	Linked GPO	Description
AllUsers	GPO1	Contains all the user accounts in the domain
AllComputers	GPO2	Contains all the computer accounts for the client computers in the domain
AllServers	GPO3	Contains all the computer accounts for Windows servers
VitualDesktops	GPO4	A new OU that will contain the computers account for Azure Virtual Desktop session hosts

Requirements:

Fabrikam Identifies the following planned changes:

* Create a single Azure subscription named Sub1 that will contain a single Azure virtual network named Vnet1.
* Replace the WAN links between the Seattle and New York offices by using Azure Virtual WAN and ExpressRoute. Both on-premises offices will be connected to Vnet1 by using ExpressRoute.
* Create three Azure file shares named newyorkfiles, seattfefiles, and companyfiles.
* Create a domain controller named dc3.corp.fabrikam,com in Vnet1.
* Deploy an Azure Virtual Desktop host pool lo Vnet1. The Azure Virtual Desktop session hosts will be hybrid Azure AD joined.
* License all servers for Microsoft Defender for servers.
* Use Azure Policy to enforce configuration management policies on the servers in Azure and on-premises.

Networking Requirements

Fabrikam identifies the following security requirements:

* Apply GP04 to the Azure Virtual Desktop session hosts. Ensure that Azure Virtual Desktop user sessions lock after being idle for 10 minutes. Users must be able to control the lockout lime manually from their client computer.
* Ensure that server administrators request approval before they can establish a Remote Desktop connection to an Azure virtual machine. If the request is approved, the connection must be established within two hours.
* Prevent user passwords from containing all or part of words that are based on the company name, such as Fab. fabrikam or fsbr! |.
* Ensure that all instances of Webapp1 use the same service account. The password of the service account must change automatically every 30 days.
* Prevent domain controllers from directly contacting hosts on the internet.

File Sharing Requirements

You need to configure the synchronization of Azure files to meet the following requirements:

* Ensure that seattlefiles syncs to FS2.

* Ensure that newyorkfiles syncs to FS1.
* Ensure that companyfiles syncs to both FS1 and FS2.

NEW QUESTION: 55
Task 9
You need to ensure that all the computers in the domain use DNSSEC to resolve names in the adatum.com zone.
Answer:
See the solution of this Task below.
Explanation:
To ensure that all computers in the domain use DNSSEC to resolve names in the adatum.com zone, you'll need to configure both the DNS servers and the client computers. Here's how you can do it:
Step 1: Sign the adatum.com Zone First, you need to sign the adatum.com DNS zone. This can be done using the DNS Manager or PowerShell. Here's a PowerShell example:
Add-DnsServerSigningKey -ZoneName "adatum.com" -CryptoAlgorithm RsaSha256 Set-DnsServerDnsSecZoneSetting -ZoneName "adatum.com" -DenialOfExistence NSEC3 -NSEC3Parameters 1,0,10,"" This will add a signing key and configure DNSSEC for the zone with NSEC3 parameters.
Step 2: Configure DNS Servers Ensure that your DNS servers are configured to support DNSSEC. This includes setting up trust anchors for the zones that you want to validate and configuring the DNS servers to provide DNSSEC validation for DNS queries.
Step 3: Configure DNS Clients For DNSSEC validation to occur on the client side, the client computers must be configured to trust the DNS server's validation process. This typically involves configuring the client's DNS settings to point to a DNS server that supports DNSSEC.
Step 4: Validate Configuration You can validate that DNSSEC is working correctly by using tools like nslookup or dig to query DNS records and check for the presence of DNSSEC signatures in the responses.
Note: The exact steps may vary depending on your environment and the version of Windows Server you are using. Ensure that you have the appropriate administrative rights to make these changes and that you test the configuration in a controlled environment before deploying it domain-wide12.
By following these steps, you should be able to ensure that all computers in your domain use DNSSEC to resolve names in the adatum.com zone.

NEW QUESTION: 56
You have an Azure subscription that contains the following resources:
* An Azure Log Analytics workspace
* An Azure Automation account
* Azure Arc.

You have an on-premises server named Server1 that is onboaraed to Azure Arc You need to manage Microsoft updates on Server! by using Azure Arc Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point
A. Add Microsoft Sentinel to the Log Analytics workspace
B. On Server1, install the Azure Monitor agent
C. From the Automation account, enable Update Management for Server1.
D. From the Virtual machines data source of the Log Analytics workspace, connect Server1.
Answer: B,C (LEAVE A REPLY)
To manage Microsoft updates on an on-premises server via Azure Arc, you must leverage the Azure Update Management solution. According to official Administering Windows Server Hybrid Core Infrastructure study guides, Azure Arc-enabled servers allow you to manage your non-Azure Windows and Linux machines as if they were native Azure resources. To facilitate update management, the infrastructure relies on a connection between the machine, an Azure Automation account, and a Log Analytics workspace.
First, you must install the Azure Monitor agent (or the Legacy Log Analytics agent, though AMA is the current standard) on Server1. This agent is responsible for collecting inventory data and the status of available updates, then sending that metadata to the Log Analytics workspace. Azure Arc simplifies this deployment by allowing the agent to be installed as an extension directly from the Azure portal.
Second, you must enable Update Management from the Automation account. The Update Management solution uses the data collected by the agent to assess the update status of the server. Once enabled, you can create update deployments to schedule the installation of updates during defined maintenance windows. This integrated approach ensures that hybrid servers-those residing on-premises but managed via Azure Arc- receive the same governance and patching consistency as cloud-native virtual machines. Option A is incorrect as Sentinel is for security orchestration, and Option D is incorrect because Arc-enabled servers are not connected through the "Virtual machines" data source used for native Azure VMs.

NEW QUESTION: 57
Your network contains a single domain Active Directory Domain Services (AD DS) forest named contoso.
com. The forest contains a single Active Directory site.
You plan to deploy a read only domain controller (RODC) to a new datacenter on a server named Server1. A user named User1 is a member of the local Administrators group on Server1.
You need to recommend a deployment plan that meets the following requirements:
Ensures that a user named User1 can perform the RODC installation on Server1 Ensures that you can control the AD DS replication schedule to the Server1 Ensures that Server1 is in a new site named RemoteSite1 Uses the principle of least privilege Which three actions should you recommend performing in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

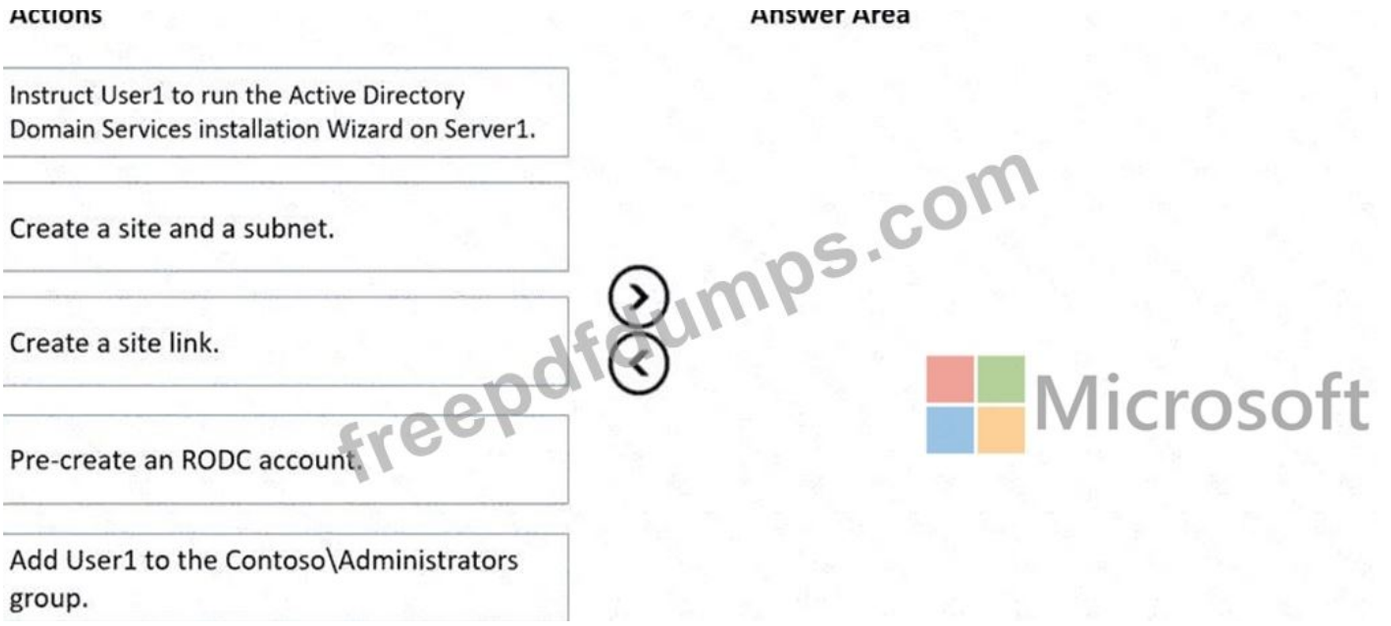

Answer:

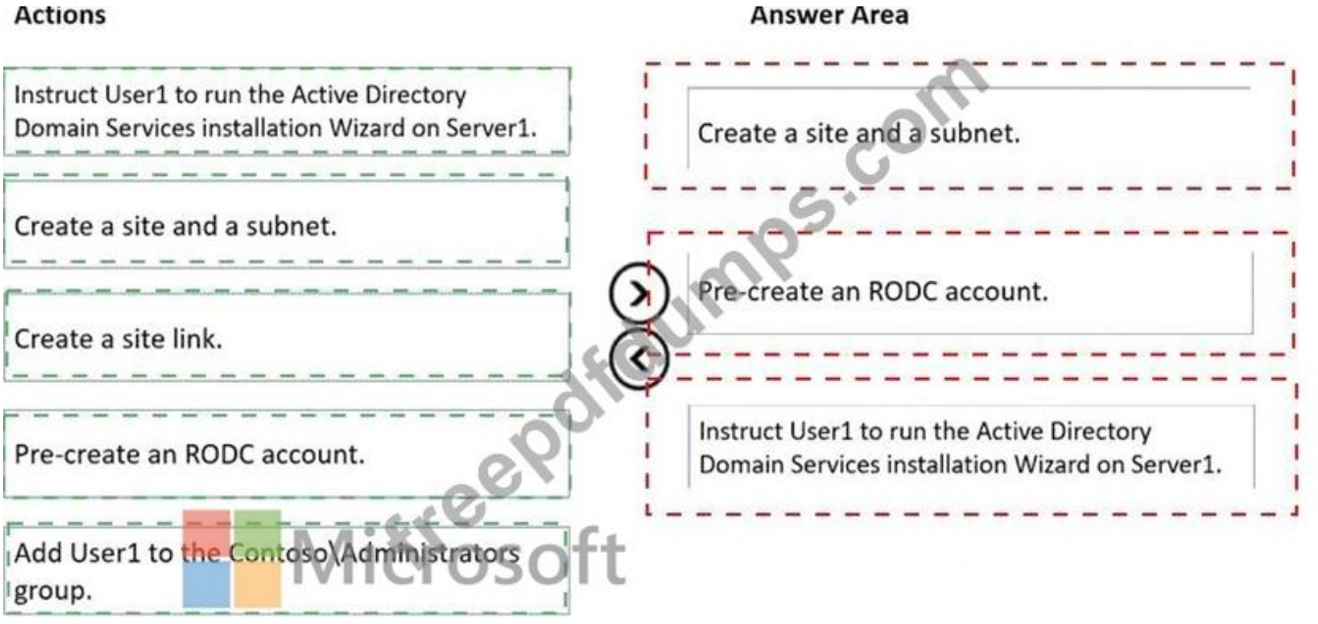

Explanation:

The Windows Server Hybrid Core Infrastructure curriculum explains that AD DS clients and DCs are site- aware and their placement is governed by site and subnet objects. To place a new RODC in a specific datacenter and ensure it belongs to a new site, you must first create the site and associate the correct IP subnet.
The guide states that "subnet-to-site mapping determines the site membership of domain controllers and clients," ensuring Server1 is in RemoteSite1 after promotion. Next, to control replication, the guidance under
"Configure sites, subnets, and site links" specifies that inter-site replication schedules and costs are configured on site links; adjusting the site link gives you granular control over when the hub site (your existing site) replicates with RemoteSite1. Finally, to meet the principle of least privilege and allow a non-Domain Admin to install an RODC, the module "Deploy and manage RODCs" describes pre-staging (pre-creating) an RODC account and delegating installation to a designated user. It notes that "an RODC can be pre-created and an installer account assigned so that a local administrator at the branch can run the AD DS installation wizard without requiring domain-wide administrative rights." With these three steps-site/subnet, site link, and pre- created RODC with delegated installer-User1 can promote Server1 to an RODC in RemoteSite1 while you retain control of replication and adhere to least privilege.

NEW QUESTION: 58
You plan to deploy a containerized application that requires .NET Core.
You need to create a container image for the application. The image must be as small as possible.
Which base image should you use?
A. Nano Server
B. Server Cote
C. Windows Server
D. Windows
Answer: A (LEAVE A REPLY)
When building Windows container images, the base image determines both compatibility and image size. The hybrid core curriculum emphasizes that Nano Server is the smallest Windows base image, intended for headless, app-only workloads such as .NET (/.NET Core) applications. .NET Core's modular, self-contained approach enables it to run on Nano Server images designed for Windows containers, yielding significantly smaller images and faster pull/start times than Server Core or Windows Server (Full) bases. Server Core includes additional components (GUI-less but still broad OS surface) and is used when you need APIs or frameworks not present in Nano. "Windows" or "Windows Server (Full)" are not container base images for modern Windows containers and would produce unnecessarily large layers. To meet the requirement "the image must be as small as possible" for a .NET Core app, the guidance is to select Nano Server as the base (for matching architecture and Windows version), then layer the .NET runtime or self-contained app on top.

This choice aligns with best practices for minimizing footprint, improving density, and reducing network transfer during CI/CD.

NEW QUESTION: 59

You have an on-premises server that runs Windows Server and contains a file share named Share1.

You have an Azure subscription that contains an Azure Files share named azshare1 and an Azure File Sync instance named Synd. Sync1 syncs Share1 with azshare1.

You need to delete Synd.

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

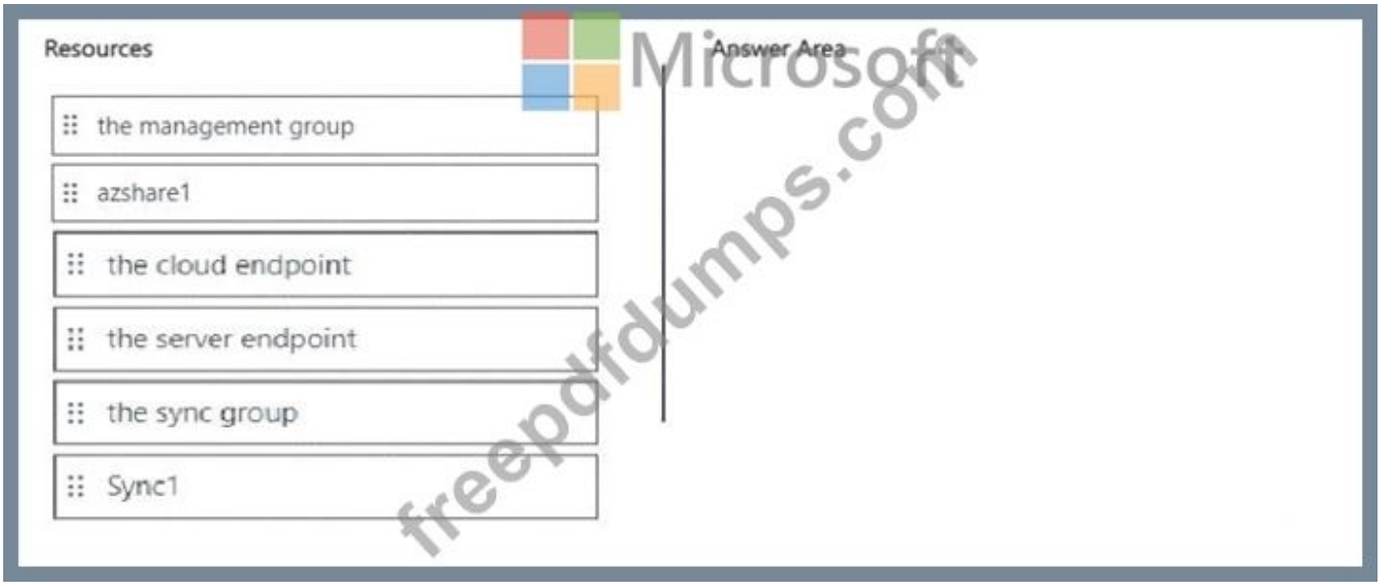

Answer:

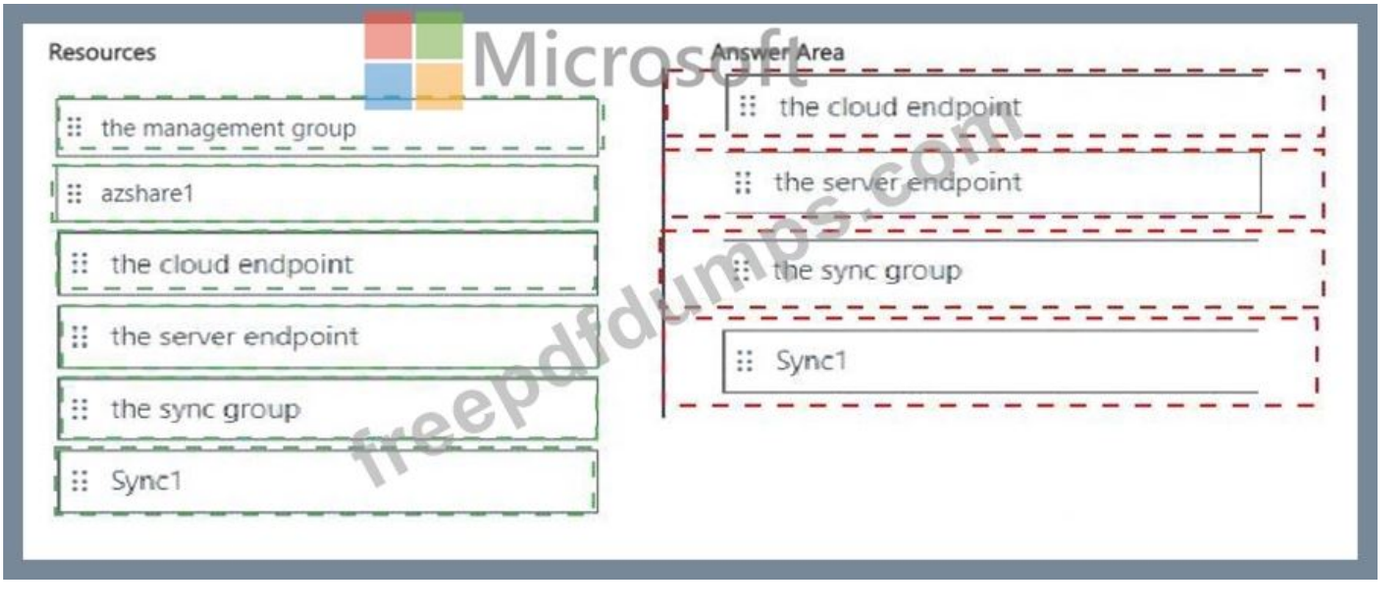

Explanation:

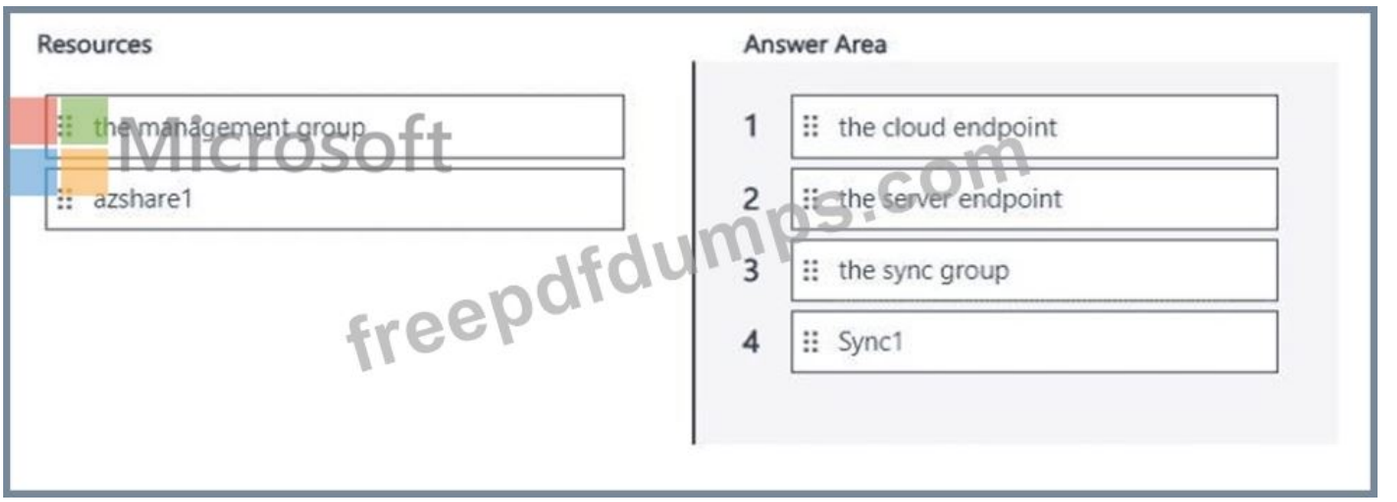

NEW QUESTION: 60
Your network contains an Active Directory Domain Services (AD DS) domain named contoso.com. The domain contains two servers named Server1 and Server2.
Server1 contains a disk named Disk2. Disk2 contains a folder named UserData. UserData is shared to the Domain Users group. Disk2 is configured for deduplication. Server1 is protected by using Azure Backup.
Server1 fails.
You connect Disk2 to Server2.
You need to ensure that you can access all the files on Disk2 as quickly as possible.
What should you do?
A. Create a storage pool.
B. Restore files from Azure Backup.
C. Install the File Server Resource Manager server role.
D. Install the Data Deduplication server role.
Answer: (SHOW ANSWER)
The Windows Server Data Deduplication section of the Administering Windows Server Hybrid Core Infrastructure course explains that deduplicated volumes require the deduplication components on any system that needs to read them: "Data Deduplication is implemented as a filter driver. To access a deduplicated volume on another server, you must install the Data Deduplication role service; otherwise files may appear as reparse points and cannot be opened." It also clarifies that "when the role is present, the system transparently rehydrates data on access without requiring a restore." In your scenario, Disk2 was deduplicated on Server1 and is now attached to Server2. To regain immediate access to all files as quickly as possible, install the Data Deduplication role service on Server2 so the filter can interpret the chunk store and metadata and make the files readable instantly. Creating a storage pool is irrelevant because the disk already exists and is readable; FSRM does not interpret dedup data; and restoring from Azure Backup would be slower and unnecessary since the original deduplicated volume is intact. Hence, the fastest, correct action is to install Data Deduplication on Server2.https://docs.microsoft.com/en-us/windows-server/storage/data-deduplication

/overview

NEW QUESTION: 61
Your network contains an Active Directory Domain Services (AD DS) domain.
You plan to use Active Directory Administrative Center to create a new user named User1.
Which two attributes are required to create User1? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. Password
B. Profile path
C. User SamAccountName logon
D. Full name
E. First name
F. User UPN logon
Answer: C,D (LEAVE A REPLY)
When creating users with Active Directory Administrative Center (ADAC), the New User workflow highlights the required attributes with indicators. The Administering Windows Server Hybrid Core Infrastructure materials note that ADAC uses a modern schema-driven form in which "Full name (CN)" and
"User UPN logon" are the minimum required identity fields to create the object in the directory. The wizard auto-generates the sAMAccountName from the UPN by default (you can edit it), but sAMAccountName isn't required to be manually entered to complete creation. Likewise, Password can be deferred depending on your provisioning pattern (for example, creating a disabled or pre-staged account or enforcing "User must change password at next logon"), and fields such as Profile path and First name are optional profile details. The guide explains that ADAC "derives the RDN from Full name" and relies on UPN as the primary modern logon attribute in Azure AD-connected/hybrid scenarios, ensuring uniqueness within the UPN suffix. Therefore, to successfully create User1 using ADAC without additional, non-mandatory properties, you must provide Full name and User UPN logon.

NEW QUESTION: 62

You have servers that have the DNS Server role installed. The servers are configured as shown in the following table.

Name	Office	Local DNS zone	IP address
Server1	Paris	contoso.com	10.1.1.1
Server2	New York	None	10.2.2.2

All the client computers in the New York office use Server2 as the DNS server.
You need to configure name resolution in the New York office to meet the following requirements:
Ensure that the client computers in New York can resolve names from contoso.com.
Ensure that Server2 forwards all DNS queries for internet hosts to 131. 107.100.200.
The solution must NOT require modifications to Server1.
Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. a forwarder
B. a conditional forwarder
C. a delegation
D. a secondary zone
E. a reverse lookup zone
Answer: A,B (LEAVE A REPLY)
Infrastructure documents: =
In the Windows Server DNS guidance for hybrid core infrastructure, Microsoft specifies that forwarders send
"all queries for names that the DNS server cannot resolve locally to a designated upstream server," while conditional forwarders target "queries for a specific DNS domain to one or more authoritative DNS servers for that domain." When a branch office DNS server must resolve a partner/remote domain hosted elsewhere in the organization and you cannot change the authoritative server, the recommended pattern is to configure a conditional forwarder on the branch server pointing to the remote authoritative server. For Internet name resolution, you configure a standard forwarder to the required recursive resolver IP.
Applied here: New York clients use Server2. To resolve contoso.com (hosted on Server1), create a conditional forwarder on Server2 for contoso.com that points to 10.1.1.1 (Server1). To meet the requirement to forward all other external lookups, configure a forwarder on Server2 to 131.107.100.200. This design avoids zone transfers or changes on Server1 and fulfills both name-resolution requirements with minimal administration.

NEW QUESTION: 63
Your on-premises network contains an Active Directory Domain Services (AD DS) domain. The domain contains the servers shown in the following table.

Name	Description
DC1	Domain naming master, PDC emulator, and RID master
DC2	Schema master and infrastructure master
RODC1	Read-only domain controller (RODC)
Server1	Microsoft Entra Connect sync server
Server2	Microsoft Entra Application Proxy connector

The domain controllers do NOT have internet connectivity.

You plan to implement Azure AD Password Protection for the domain.

You need to deploy Azure AD Password Protection agents. The solution must meet the following requirements:

* All Azure AD Password Protection policies must be enforced.
* Agent updates must be applied automatically.
* Administrative effort must be minimized.

What should you do? To answer select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

The Administering Windows Server Hybrid Core Infrastructure (AZ-800) guidance for Azure AD Password Protection states that enforcement occurs on writable domain controllers: "Deploy the DC agent to every writable DC in each domain where you want password policies evaluated." It

further clarifies: "Do not install the DC agent on RODCs; RODCs don't perform password set/change operations." Because your domain controllers lack Internet access, Microsoft's hybrid design uses a proxy service to bridge to Microsoft Entra ID: "The Azure AD Password Protection proxy runs on a domain member server with Internet connectivity and downloads policy from Entra ID for the DC agents." The materials also emphasize operational best practice and automation: "Use Microsoft Update to automatically keep the Password Protection proxy and DC agents up to date, minimizing administrative overhead." Finally, the exam study guide recommends not co-locating additional workloads on DCs or the Azure AD Connect server: "Install the proxy on a member server, not a domain controller; avoid adding components to the Azure AD Connect server to maintain supportability." Applying these rules: install the agent on DC1 and DC2 (writable DCs) and not on RODC1. Place the proxy on an Internet-connected member server-Server2 (already an Application Proxy connector)-to meet enforcement, automatic updates, and minimal administrative effort.

NEW QUESTION: 64

Task 3

You need to configure SRV1 as a DNS server. SRV1 must be able resolve names from the contoso.com domain by using DC1. All other names must be resolved by using the root hint servers.

Answer:

See the solution of this Task below.

Explanation:

One possible solution to configure SRV1 as a DNS server that can resolve names from the contoso.com domain by using DC1 and all other names by using the root hint servers is to use conditional forwarding.

Conditional forwarding allows a DNS server to forward queries for a specific domain name to another DNS server, while using the normal forwarding or root hint servers for other queries. Here are the steps to configure conditional forwarding on SRV1:

* On SRV1, open DNS Manager from the Administrative Tools menu or by typing dnsmgmt.msc in the Run box.
* In the left pane, right-click on Conditional Forwarders and select New Conditional Forwarder.
* In the New Conditional Forwarder dialog box, enter contoso.com as the DNS Domain name.
* In the IP addresses of the master servers box, enter the IP address of DC1, which is the DNS server for the contoso.com domain. You can also click on Resolve to verify the name resolution of DC1.
* Optionally, you can check the box Store this conditional forwarder in Active Directory, and replicate it as follows if you want to store and replicate the conditional forwarder in AD DS. You can also select the replication scope from the drop-down list.
* Click OK to create the conditional forwarder.

Now, SRV1 will forward any queries for the contoso.com domain to DC1, and use the root hint servers for any other queries. You can test the name resolution by using the nslookup command

on SRV1 or another computer that uses SRV1 as its DNS server. For example, you can run the following commands:
nslookup www.contoso.com
nslookup www.microsoft.com
The first command should return the IP address of www.contoso.com from DC1, and the second command should return the IP address of www.microsoft.com from a root hint server.

NEW QUESTION: 65
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 You need to create a Docker network that supports the deployment of the application. Which type of network should you create?
A. transparent
B. l2bridge
C. NAT
D. l2tunnel
Answer: (SHOW ANSWER)
In the context of Windows Server Hybrid Core Infrastructure and container networking, choosing the correct network driver is critical for application deployment. According to official documentation, the l2bridge (Layer 2 Bridge) network mode is used when container hosts are connected to the same IP subnet. In this configuration, each container is assigned an IP address from the same prefix as the container host. All container traffic is bridged to the physical network through an external Hyper-V Virtual Switch. Because the containers share the same underlying network infrastructure as the host, they are visible to the rest of the physical network without requiring Network Address Translation (NAT).
The documentation specifies that for multi-node clusters or deployments where containers must be directly reachable on the physical network via their own IP addresses, l2bridge is the standard choice. This differs from NAT, which uses a private internal IP range and translates traffic through the host's IP, and Transparent mode, which is often used for individual hosts where the container is directly connected to the physical network but can have complexities in virtualized environments. l2tunnel is specifically used for Microsoft Cloud Stack (Azure Stack HCI) and SDN scenarios, typically involving encapsulation, which is not the standard requirement for a general multi-container application deployment on a single server unless specified.
Therefore, for a high-performance, direct-access network that bridges traffic at Layer 2, l2bridge is the verified architectural choice for Windows containers.

NEW QUESTION: 66
You have a server named Server1 that has the Hyper-V server role installed. Server1 hosts the virtual machines shown in the following exhibit.

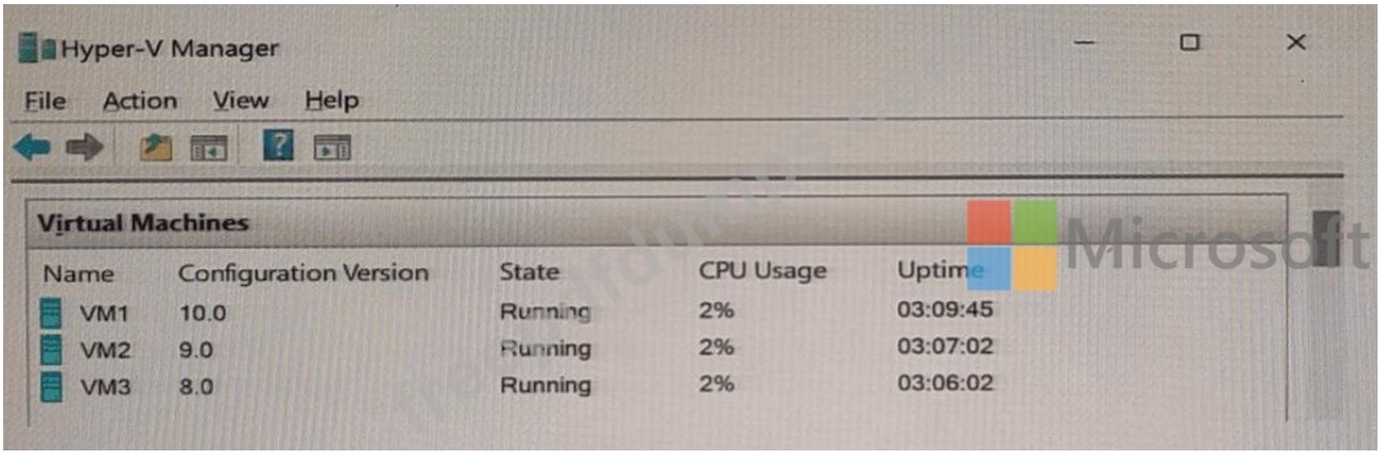

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

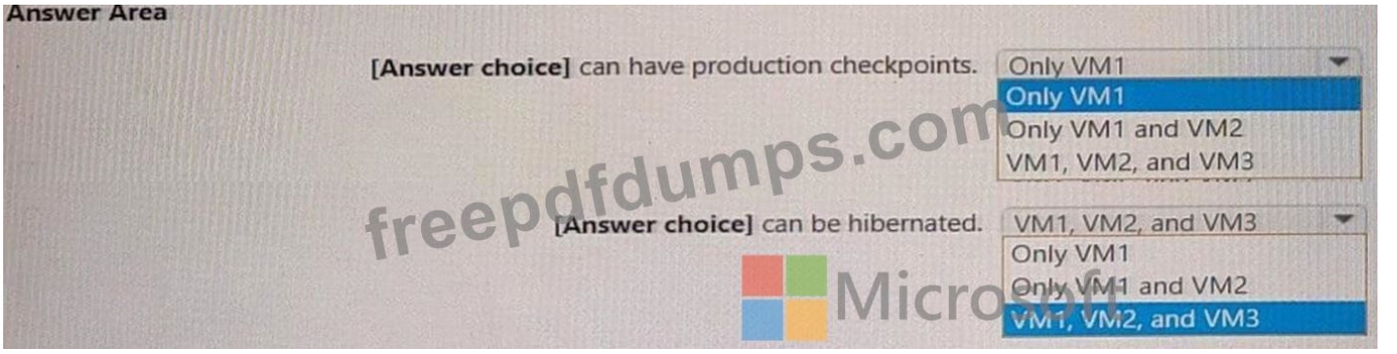

Answer:

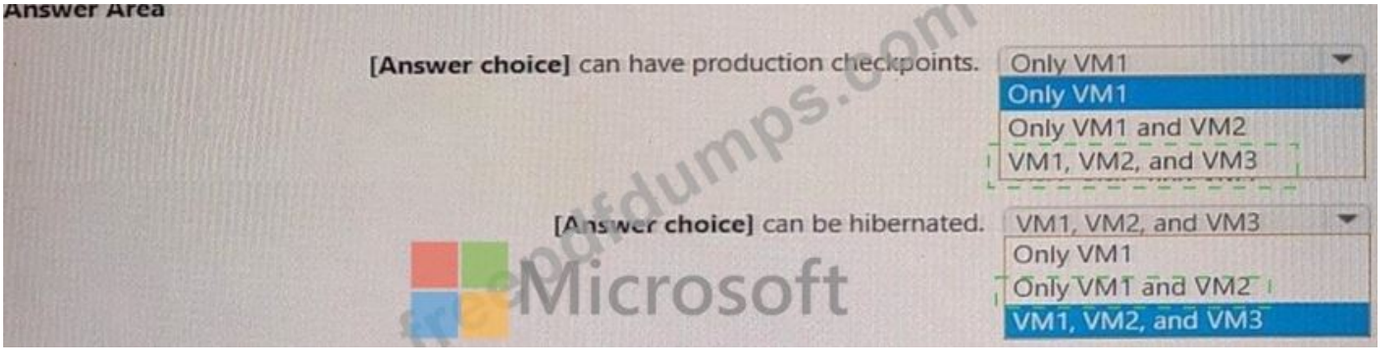

Explanation:

[Answer choice] can have production checkpoints. = VM1, VM2, and VM3

[Answer choice] can be hibernated. = Only VM1 and VM2

NEW QUESTION: 67

You have an on-premises server named Server1 that runs Windows Server and has internet connectivity.

You have an Azure subscription.

You need to monitor Server1 by using Azure Monitor.

Which resources should you create in the subscription, and what should you install on Server1?

To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

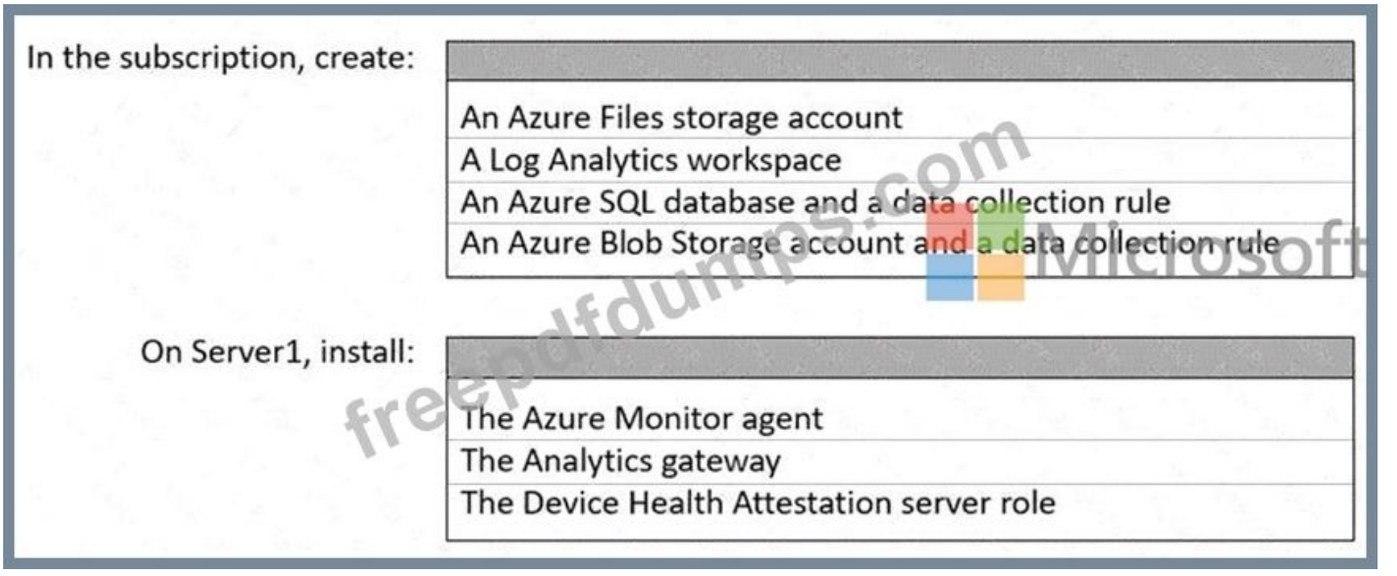

Answer:

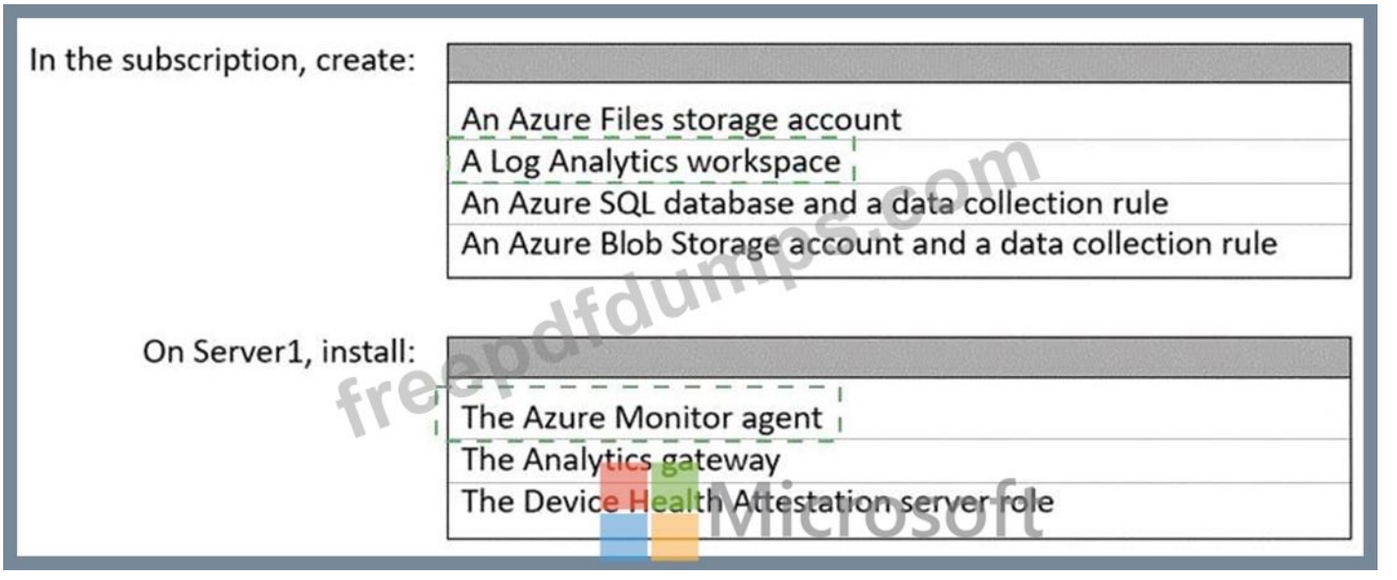

Explanation:

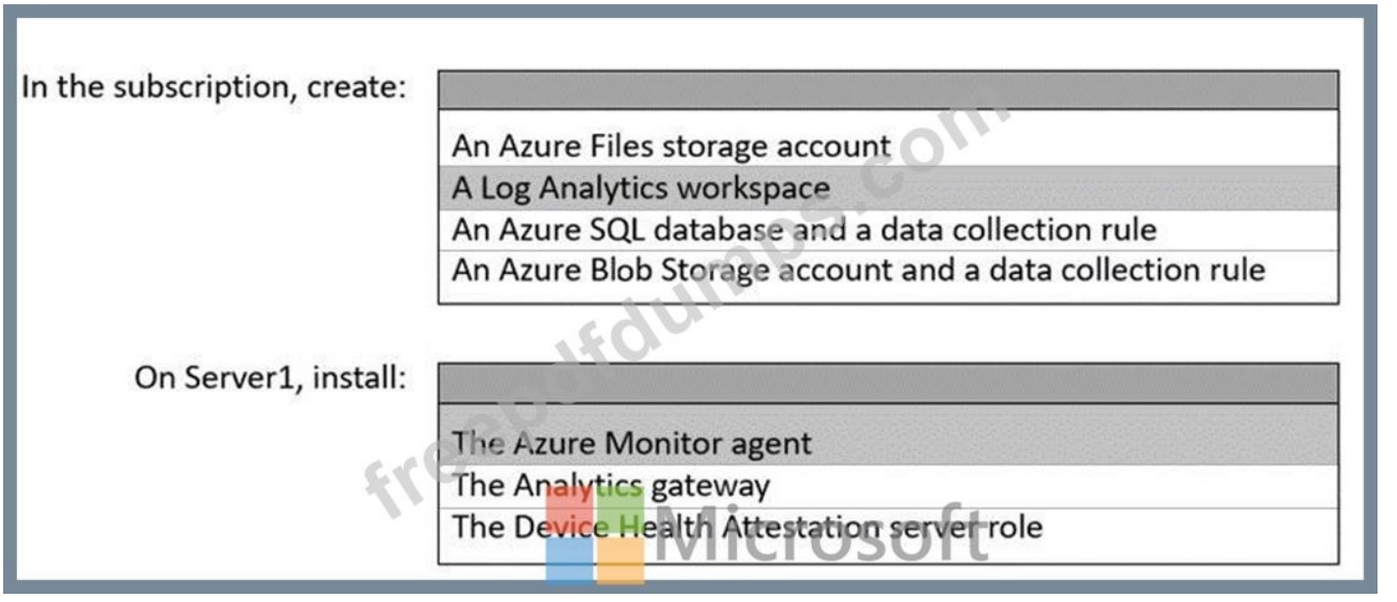

The Administering Windows Server Hybrid Core Infrastructure learning path explains that onboarding Windows Server machines to Azure Monitor hinges on a Log Analytics workspace as

the data store for logs and metrics, and an agent on the machine to collect and send data. The guide states that "Azure Monitor collects telemetry from Windows Server machines into a Log Analytics workspace," and that for current deployments you should "install the Azure Monitor agent (AMA) on each computer and associate it with a workspace." It further notes that "the AMA replaces the legacy Log Analytics agent" and "uses data collection rules (DCRs) to define which data to collect and to which Log Analytics workspace to send it." Storage accounts, SQL databases, or analytics gateways are not required to onboard a standalone Windows Server to Azure Monitor. Likewise, Device Health Attestation is unrelated to Azure Monitor onboarding.
In short, to monitor an on-premises Windows Server from Azure you create a Log Analytics workspace in the subscription and install the Azure Monitor agent on the server. (In practice, you then configure a DCR to route data from the AMA to the workspace, but no additional Azure resources like SQL or Blob/Azure Files are needed for basic monitoring.)

NEW QUESTION: 68
You have the servers shown in the following tab

Name	Role
Server1	Hyper-V
Server2	Hyper-V
Server3	DHCP Server

Served contains a virtual machine named VM1 that runs Windows Server. Server1 has an external switch named Switch"1. VM1 is connected to Switch1.
You provision containers on VM1.
You need to configure networking for VM1. The solution must meet the following requirements:
* Ensure that Server3 automatically assigns IP addresses to the containers.
* Ensure that the containers can communicate with Server2.
What should you do? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

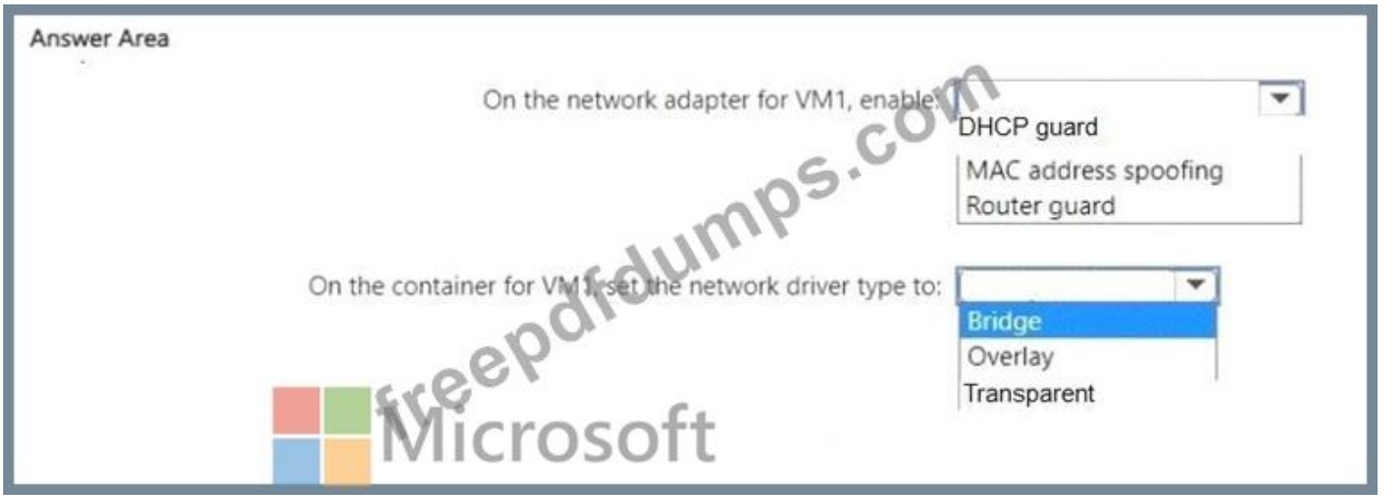

Answer:

Explanation:

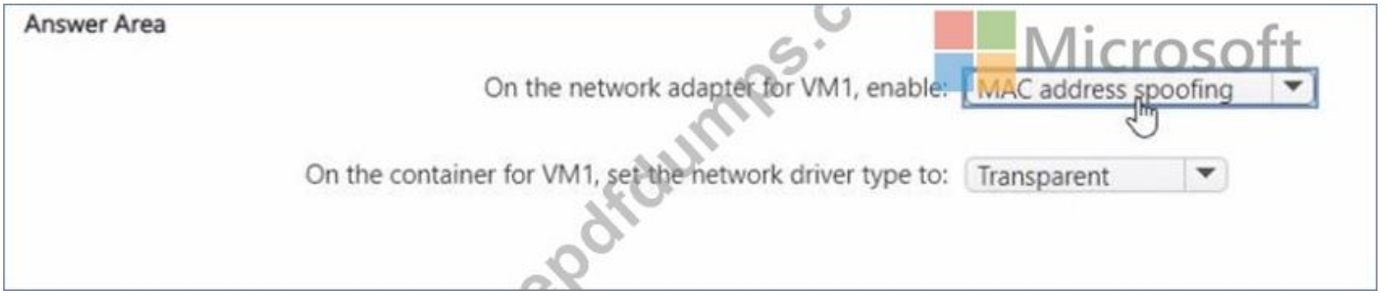

NEW QUESTION: 69

You have four testing devices that are configured with static IP addresses as shown in the following table.

Name	IP address
TestDevice1	192.168.16.242
TestDevice2	192.168.16.243
TestDevice3	192.168.16.244
TestDevice4	192.168.16.245

The test devices are turned on once a month.

You need to prevent Server1 from assigning the IP addresses allocated to the test devices to other devices when the test devices are offline. The solution must minimize administrative effort.

What should you do?

A. Create an exclusion range.

B. Create reservations.

C. Configure the Scope options.

D. Create a policy.

Answer: A (LEAVE A REPLY)

The DHCP planning and configuration section in Administering Windows Server Hybrid Core Infrastructure states: "Use exclusion ranges to prevent the DHCP Server from leasing specific addresses that are in use by statically assigned hosts." Reservations are intended for DHCP clients that must always receive the same IP from the scope; they do not apply to devices configured with static IPs and would add unnecessary administration. Because the four test devices are statically addressed and are powered on only monthly, the simplest way to ensure

the DHCP server never leases their addresses is to exclude those exact IP addresses from the scope. Scope options and policies do not reserve or block individual addresses from lease assignment. Therefore, to prevent conflicts with minimal effort, configure exclusion entries for 192.168.16.242-245 in the relevant scope.

NEW QUESTION: 70
You have an Active Directory Domain Services (AD DS) domain. The domain contains a member server named Server1 that runs Windows Server.
You need to ensure that you can manage password policies for the domain from Serve1.
Which command should you run first on Server1?
A. Install-Windows Feature RSAT-AO-PowerShell
B. Install-Windows Feature 6PHC
C. Install-Windows Feature RSAT-AD-Tool$
D. Install-windows Feature RSAT-AWIMS
Answer: C (LEAVE A REPLY)
In Windows Server hybrid environments, domain password policy and fine-grained password policy (PSO) administration is performed with the Active Directory administration tools. The AZ-800 study domain states that to "manage AD DS, including users, groups, OUs, Group Policy, and fine-grained password policies from a member server, you must first install the Remote Server Administration Tools (RSAT) for Active Directory." The RSAT AD DS package (RSAT-AD-Tools) installs Active Directory Users and Computers (ADUC), Active Directory Administrative Center (ADAC) (which includes the Password Settings container UI), and Windows PowerShell for Active Directory (the ActiveDirectory module). The guidance further clarifies: "ADAC exposes Password Settings Objects (PSOs) and lets administrators create, link, and manage fine-grained password and lockout policies without signing in to a domain controller." Therefore, on Server1 the first required command is: Install-WindowsFeature RSAT-AD-Tools. Options for only PowerShell, DHCP, or WINS features do not provide ADAC/ADUC and cannot manage PSOs or domain password policies from a member server.

NEW QUESTION: 71
You need to meet the technical requirements for Server4.
Which cmdlets should you run on Server1 and Server4? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

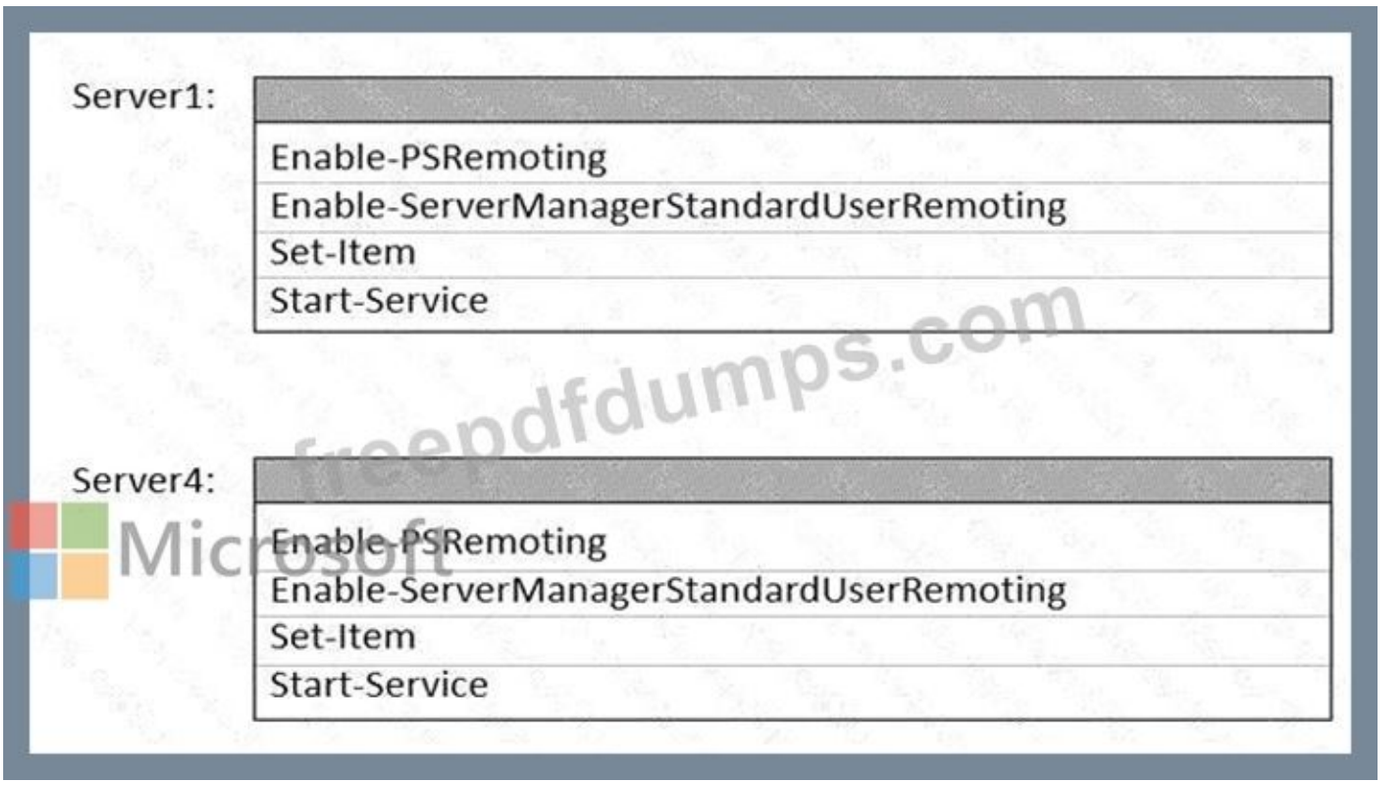

Answer:

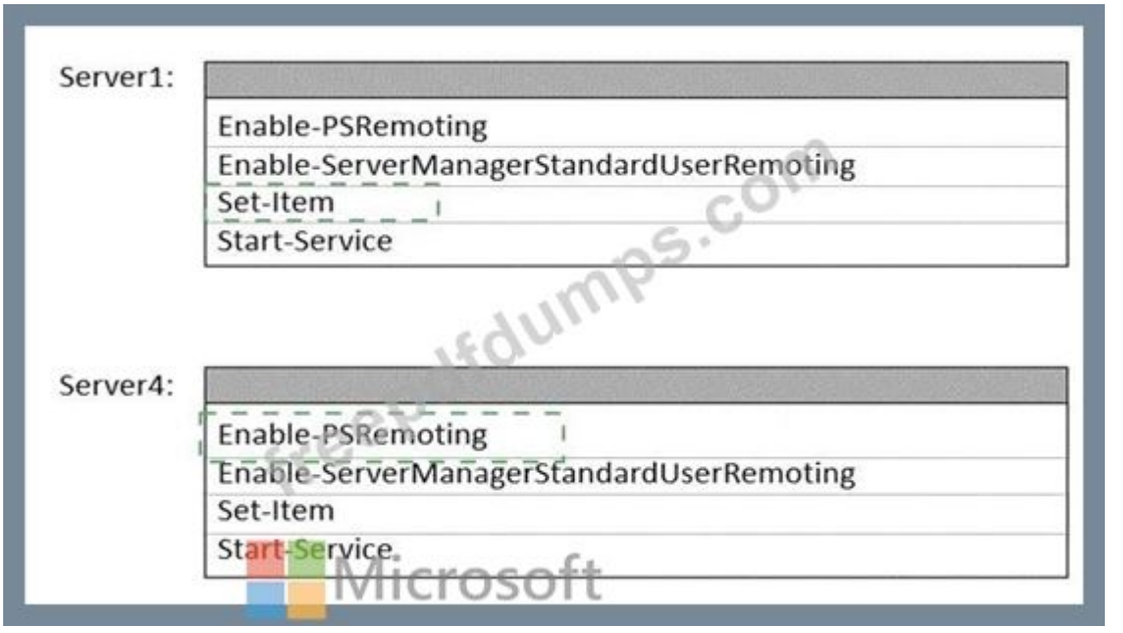

Explanation:

Server1 - Set-Item

Server4 - Enable-PSRemoting

NEW QUESTION: 72

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

Name	Type
Server1	Physical server
VM2	Hyper-V virtual machine

You have an Azure subscription that contains a virtual machine named VMV You need to ensure that you can manage all the servers by using Azure Arc. The solution must minimize administrative effort.
On which servers should you install the Azure Connected Machine agent?
A. Server1 only
B. VM1 only
C. VM2only
D. VM1 and VM2 only
E. Server1 and VM2 only
F. Server1, VM1, and VM2
Answer: E (LEAVE A REPLY)
Azure Arc-enabled servers use the Azure Connected Machine agent to onboard non-Azure machines (on- premises or other clouds) so they appear in Azure for inventory, policy, Update Management, Defender, etc.
The AZ-800 study guide emphasizes: "Install the Connected Machine agent on Windows or Linux servers that are outside of Azure. Native Azure VMs are already Azure resources and can be managed without Arc; onboarding Azure VMs to Arc is optional and not required for basic management." In the scenario, Server1 (a physical on-premises server) and VM2 (a Hyper-V VM on-premises) are non-Azure; both require the Connected Machine agent to be projected into Azure via Arc. VM1 is already an Azure VM and can be managed through Azure natively (Azure VM resource model, extensions, policies) without installing the Arc agent. Because the requirement is to manage all servers via Azure Arc while minimizing administrative effort, you install the agent only where it's needed-on Server1 and VM2. This provides Arc governance and management for on-prem systems without redundant configuration on the Azure VM.

NEW QUESTION: 73
You haw? a server named Host1 that has the Hyper-V server role installed. Host! hosts a virtual machine named VM1.
You have a management server named Server! that runs Windows Server. You remotely manage Host1 from Server1 by using Hyper-V Manager.
You need to ensure that you can access a USB hard drive connected to Server1 when you connect to VM1 by using Virtual Machine Connection.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. From the Hyper-V Settings of Host1, select Allow enhanced session mode
B. From Disk Management on Host1. attach a virtual hard disk.
C. From Virtual Machine Connection, switch to a basic session.
D. From Virtual Machine Connection select Show Options and then select the USB hard drive.
E. From Disk Management on Host1, select Rescan Disks
Answer: A,D (LEAVE A REPLY)

In Hyper-V, access to local devices such as a USB hard drive connected to the management computer (Server1) from within a VM session (VM1) is provided only when using Enhanced Session Mode through the Virtual Machine Connection (vmconnect) client. The AZ-800/"Administering Windows Server Hybrid Core Infrastructure" materials emphasize that Enhanced Session Mode "enables redirection of local resources (audio, clipboard, printers, drives, and supported Plug and Play devices) to a guest when connecting with Virtual Machine Connection." To use it, you must first enable it on the host and then select the device in vmconnect before establishing the session.
Accordingly, you should:
* Enable Enhanced Session on the host: In Hyper-V Manager # Host1 # Hyper-V Settings # Enhanced Session Mode Policy, select Allow enhanced session mode (and optionally Use enhanced session mode under User settings). This satisfies option A.
* Choose the USB drive in the vmconnect UI: When launching the connection to VM1, click Show Options # Local Resources # More..., and select Drives or the specific USB device so it is redirected into the guest. This matches option D.
Options B and E relate to host Disk Management operations and do not provide guest redirection of a USB attached to Server1. Option C ("switch to a basic session") explicitly disables the RDP-based redirection channel that Enhanced Session Mode relies upon, preventing USB/drive passthrough. Therefore, the correct pair is A and D.

NEW QUESTION: 74
You need to meet the technical requirements for VM1.
Which cmdlet should you run first? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

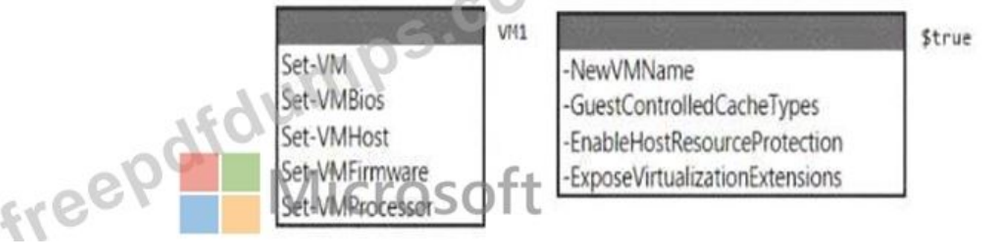

Answer:

Answer Area

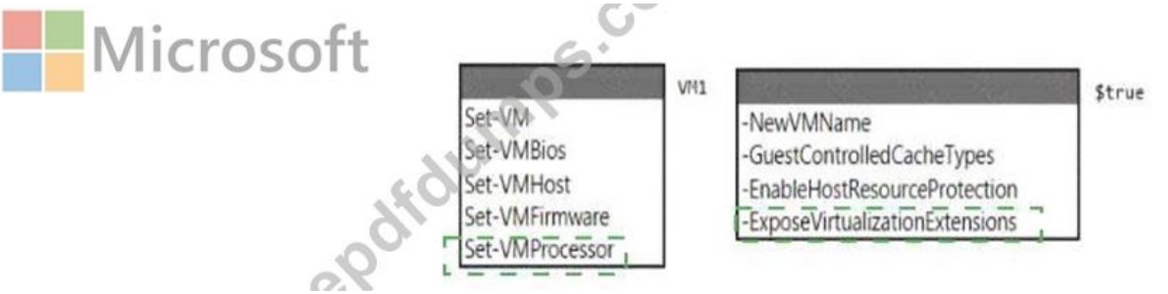

Explanation:

In the Administering Windows Server Hybrid Core Infrastructure objectives for managing Hyper-V, enabling nested virtualization is the required step when you must "run virtual machines inside a virtual machine." The referenced guidance states that Hyper-V on a VM is supported only when the host exposes hardware virtualization features to the guest. The prescriptive step is: "Turn off the VM and run Set-VMProcessor - VMName <VMName> -ExposeVirtualizationExtensions $true to enable nested virtualization." The module further notes that this action "passes through Intel VT-x/AMD-V to the guest so the guest OS can install the Hyper-V role and create VMs." It also clarifies that "the setting is applied on the parent host for the target VM and requires the VM to be powered off before the change is committed." Because the technical requirement says "Ensure that you can run virtual machines on VM1", VM1 must be able to host Hyper-V while itself running as a VM on Server2. The first and essential cmdlet is therefore Set- VMProcessor with the -ExposeVirtualizationExtensions switch set to $true against VM1. Other optional settings (for example, MAC spoofing on the vNIC or static memory) may be configured later if needed, but exposing virtualization extensions is the enabling prerequisite that satisfies the requirement.

NEW QUESTION: 75

Your network contains an Active Directory Domain Services (AD DS) domain named contoso.com. The domain contains a DNS server named Server1. Server1 hosts a DNS zone named fabrikam.com that was signed by DNSSEC.

You need to ensure that all the member servers in the domain perform DNSSEC validation for the fabrikam.

com namespace.

What should you do?

A. On Served, run the Add-DnsServerTrustAnchor cmdlet.

B. On each member server, run the Add-DnsServerTrustAnchor cmdlet.

C. From a Group Policy Object (GPO). add a rule to the Name Resolution Policy Table (NRPT).

D. From a Group Policy Object (GPO). modify the Network List Manager policies.

Answer: C (LEAVE A REPLY)

The hybrid core curriculum explains that DNSSEC validation by Windows clients is controlled through the Name Resolution Policy Table (NRPT), deployable via Group Policy. The NRPT lets administrators " require DNSSEC for specific namespaces and configure how clients validate responses." While trust anchors (Add-DnsServerTrustAnchor) are used by DNS servers performing validation, member servers acting as DNS clients rely on NRPT rules to demand

DNSSEC-validated answers from their resolvers for named namespaces (e.g., fabrikam.com). The guidance emphasizes: to "enforce DNSSEC validation on domain-joined clients for a given suffix, create a GPO-based NRPT rule that requires DNSSEC," ensuring unsigned or invalid answers are rejected. Therefore, to make all member servers validate DNSSEC for fabrikam.com, deploy a GPO NRPT rule targeting that namespace. Adding trust anchors on Server1 or on each member server is unnecessary (and in the latter case, inapplicable unless they run the DNS Server role).

NEW QUESTION: 76

Your on-premises network contains an Active Directory Domain Services (AD DS) domain.

You plan to sync the domain with a Microsoft Entra tenant by using Microsoft Entra Connect cloud sync.

You need to meet the following requirements:

* Install the software required to sync the domain and Microsoft Entra ID.
* Enable password hash synchronization.

What should you install, and what should you use to enable password hash synchronization? To answer, select the appropriate options in the answer area.

Answer Area

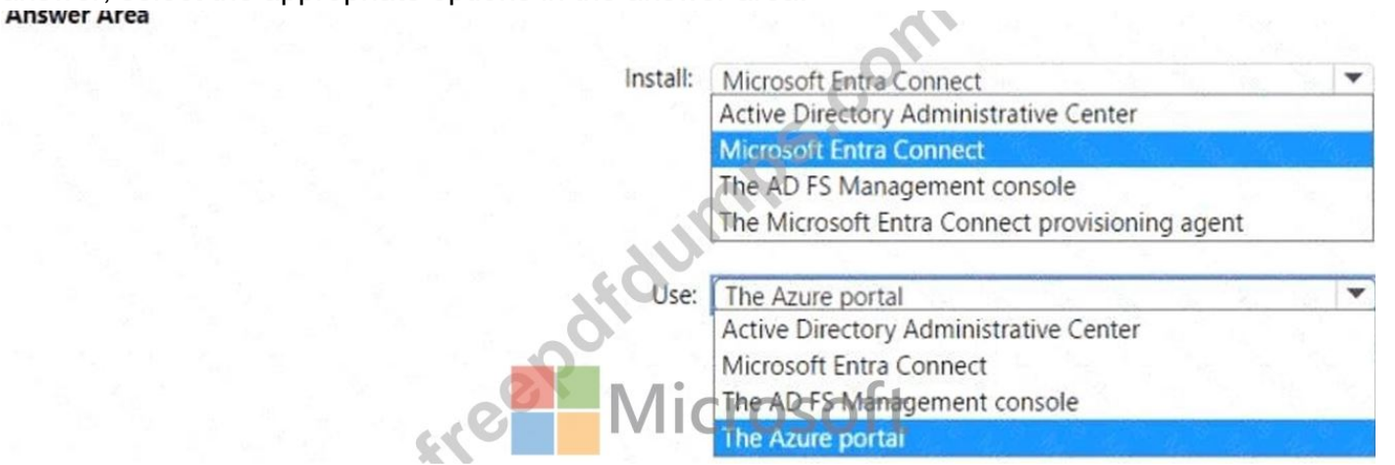

Answer:

Answer Area

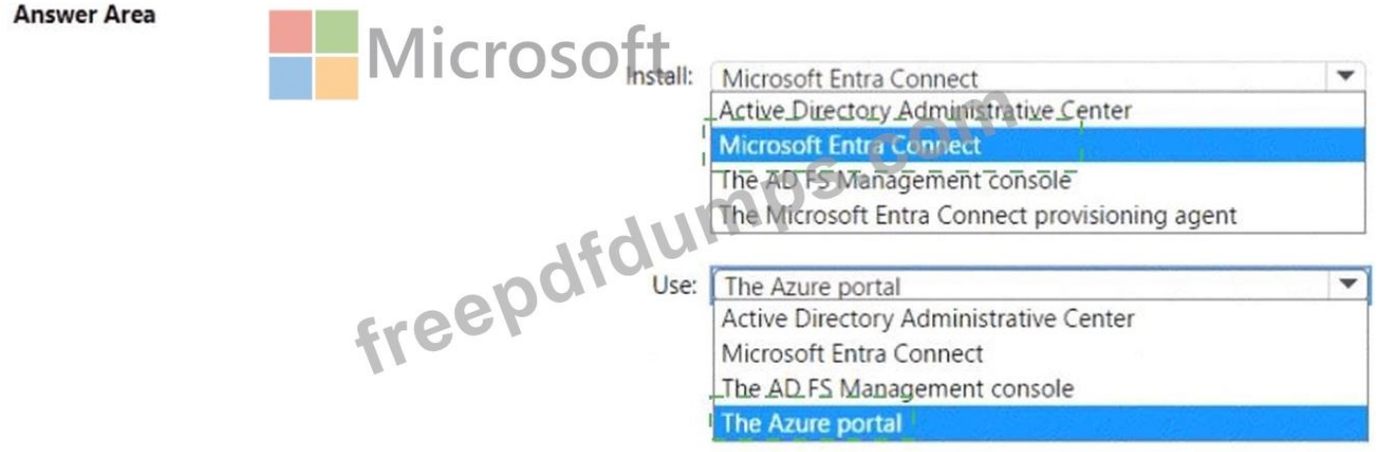

Explanation:

In a cloud-sync deployment, Microsoft's hybrid identity guidance for Administering Windows Server Hybrid Core Infrastructure specifies that you do not install the full Microsoft Entra (Azure AD) Connect server.

Instead, cloud sync "uses a lightweight agent-called the Microsoft Entra Connect provisioning agent- installed on one or more Windows Server computers that can reach your AD DS domain controllers." The agent handles directory read operations and securely communicates with the cloud provisioning service.

The study material further states that all configuration-including creating sync configurations, scoping filters, and enabling features such as Password Hash Synchronization (PHS)-is performed in the Microsoft Entra admin experience: "Cloud sync is configured in the Azure portal; you create a cloud sync configuration, select the on-premises AD forest, and enable password hash synchronization so password hashes are synchronized to Microsoft Entra ID." Because you are using Microsoft Entra Connect cloud sync, the correct software to install on-premises is the Microsoft Entra Connect provisioning agent (not the full Microsoft Entra Connect server or ADFS tools). And to enable PHS for cloud sync, you use the Azure portal (Microsoft Entra admin center) to turn on Password hash synchronization within the cloud sync configuration. This meets both requirements:

installing the proper agent for cloud sync and enabling PHS centrally in the portal.

NEW QUESTION: 77

You have an Active Directory domain that contains a file server named Server1. Server1 runs Windows Server and includes the file shares shown in the following table.

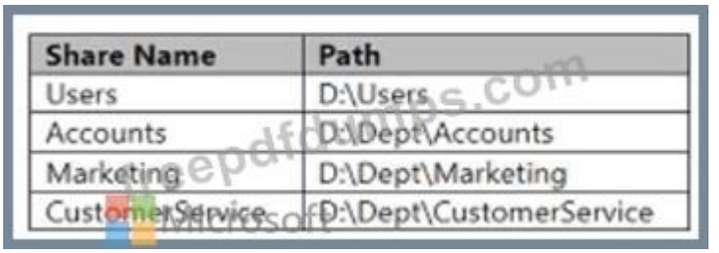

Share Name	Path
Users	D:\Users
Accounts	D:\Dept\Accounts
Marketing	D:\Dept\Marketing
CustomerService	D:\Dept\CustomerService

When users login to the network they receive the following network drive mappings.
* H: maps to Wserver1\users\%UserName%
* G: maps to \\server1\%Department%
You need to limit the amount of space consumed by user's on Server!. The solution must meet the following requirements:
* Prevent users using more than 5GB of space on their H: drive
* Prevent Accounts department users from using more than 10GB of space on the G: drive
* Prevent Marketing department users from using more than 15GB of space on the G: drive
* Prevent Customer Service department users from using more than 2GB of space on the G: drive
* Minimize administrative effort
What should you use?
A. File Server Resource Manager (FSRM) quotas
B. Storage tiering
C. NTFS Disk quotas
D. Group Policy Preferences
Answer: (SHOW ANSWER)
The Windows Server file services module specifies that FSRM quotas allow folder-level and auto-apply templates to enforce different storage limits per path, user home folder, or department share with minimal administration. The guide explains: "Use Quota Templates and Auto Apply to automatically create a separate quota for each existing and new subfolder" (e.g., D:\Users %UserName%). It also notes that quota policies can be assigned with different sizes per department folder (e.g., D:\Dept\Accounts, Marketing, CustomerService), providing email alerts and hard/soft limits. In contrast, NTFS disk quotas operate at the volume level only and cannot set distinct limits per subfolder or department; they would enforce one policy across the entire volume, which doesn't meet the requirement for multiple different limits. Storage tiering is unrelated to capacity control per user/department, and Group Policy Preferences cannot enforce storage limits. With FSRM, you would: (1) apply an auto-apply per-user hard quota of 5 GB to D:\Users, and (2) configure hard quotas of 10 GB, 15 GB, and 2 GB for the respective department folders under D:\Dept, satisfying all requirements while minimizing administrative effort.

NEW QUESTION: 78
DC1 fails.
You need to meet the technical requirements for the schema master.
Yourunntdsutil.exe.
Which five commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order?

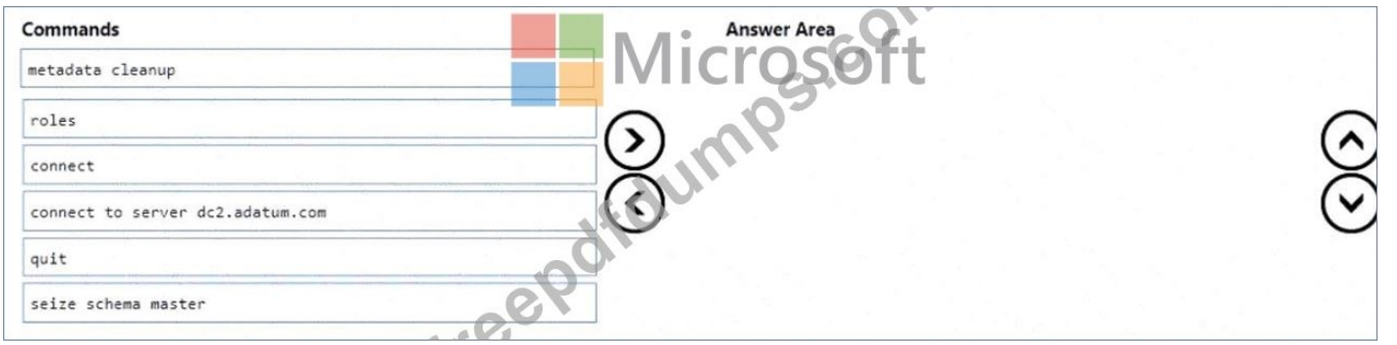

Answer:

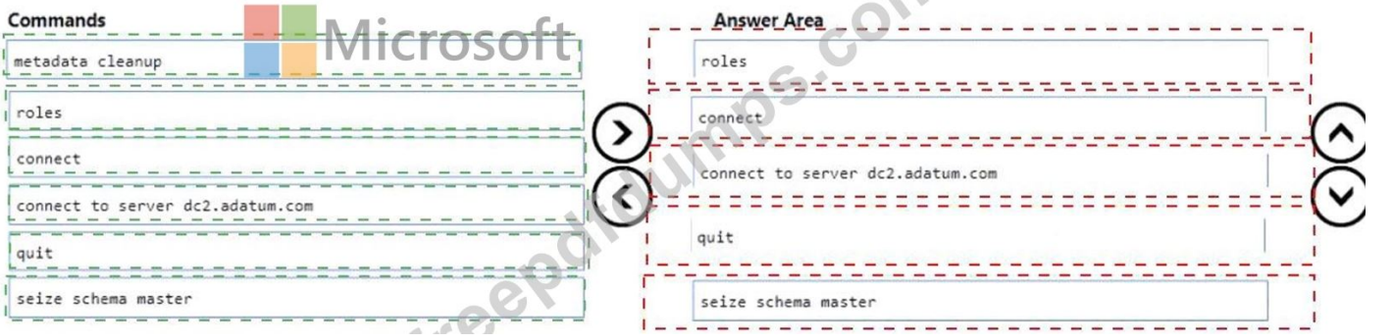

Explanation:

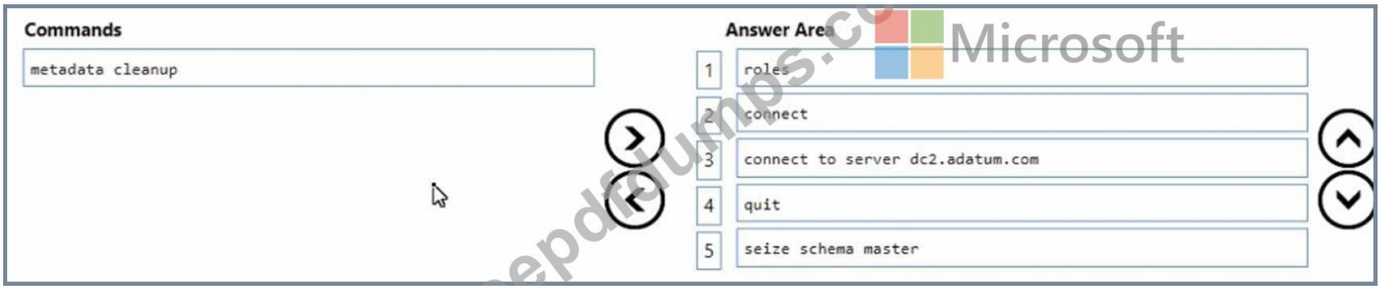

NEW QUESTION: 79

Your network contains two Active Directory Domain Services (AD DS) forests named contoso.com and fabrikam.com. Contoso.com contains three child domains named amer.contoso.com, apac.contoso.com, and emea.contoso.com. Fabrikam.com contains a child domain named apac.fabrikam.com. A bidirectional forest trust exists between contoso.com and fabrikam.com.

You need to provide users in the contoso.com forest with access to the resources in the fabrikam.com forest.

The solution must meet the following requirements:

* Users in contoso.com must only be added directly to groups in the contoso.com forest.
* Permissions to access the resources in fabrikam.com must only be granted directly to groups in the fabrikam.

com forest.

* The number of groups must be minimized.

Which type of groups should you use to organize the users and to assign permissions? To answer, drag the appropriate group types to the correct requirements. Each group type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

The Windows Server hybrid core administration guidance describes the classic cross-forest/group nesting model as AGUDLP: "Accounts # Global (or Universal) # Domain Local # Permissions." For resource access that spans forests, the documentation stresses two points: (1) Put user accounts into groups that exist only in their own forest, and (2) Grant ACL permissions only to groups that exist in the resource forest. It further explains that Universal groups are forest-wide and "can contain users and global groups from any domain in the same forest and can be used across forest trusts," which minimizes the number of groups when users come from multiple child domains. By contrast, Domain Global groups are limited to membership from a single domain, which would require separate globals per child domain and thus more groups. For assigning permissions at the resource side, the guidance states that Domain Local groups are intended to "receive permissions on resources within their own domain and can contain members from trusted forests," which satisfies the requirement that permissions in fabrikam.com be granted only to groups in that forest. Therefore, to meet all constraints and minimize group count: place contoso.com users in a Universal group (in contoso.
com) and add that group to a Domain Local group (in fabrikam.com) that is granted the required permissions.

NEW QUESTION: 80

You haw an Azure virtual machine named VM1 that runs Windows Server

You need to configure the management of VM1 to meet the following requirements:

* Require administrators to request access to VM1 before establishing a Remote Desktop connection.
* Limit access to VM1 from specific source IP addresses.
* Limit access to VMI to a specific management port

What should you configure?

A. a network security group (NSG)

B. Azure Active Directory (Azure AD) Privileged identity Management (PIM)

C. Azure Front Door

D. Microsoft Defender for Cloud

Answer: D (LEAVE A REPLY)

Reference:

The exam objectives around managing Windows Server IaaS VMs in Azure highlight Just-In-Time (JIT) VM access, a capability surfaced through Microsoft Defender for Cloud, to harden management ports. JIT enforces that administrators request time-bound access before RDP/SSH is opened, and, when approved, Defender for Cloud programmatically updates the NSG to (1) open only the specified management port, (2) for a limited time window, and (3) restricted to approved source IP addresses. This exactly satisfies: "require administrators to request access before RDP," "limit access from specific source IPs," and "limit to a specific management port." While NSGs themselves can restrict ports and source IPs, they cannot implement the approval workflow and automatic time-bound opening/closure. Azure AD PIM manages role elevation, not per-VM port exposure. Azure Front Door is for web/app delivery and does not govern RDP ingress.

Therefore, configuring JIT VM access in Microsoft Defender for Cloud is the prescribed solution in the hybrid core guidance for secure, just-in-time RDP to Windows Server VMs.

NEW QUESTION: 81

You have 10 on-premises servers that run Windows Server.

You plan to use Azure Network Adapter to connect the servers to the resources in Azure.

Which prerequisites do you require on-premises and in Azure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

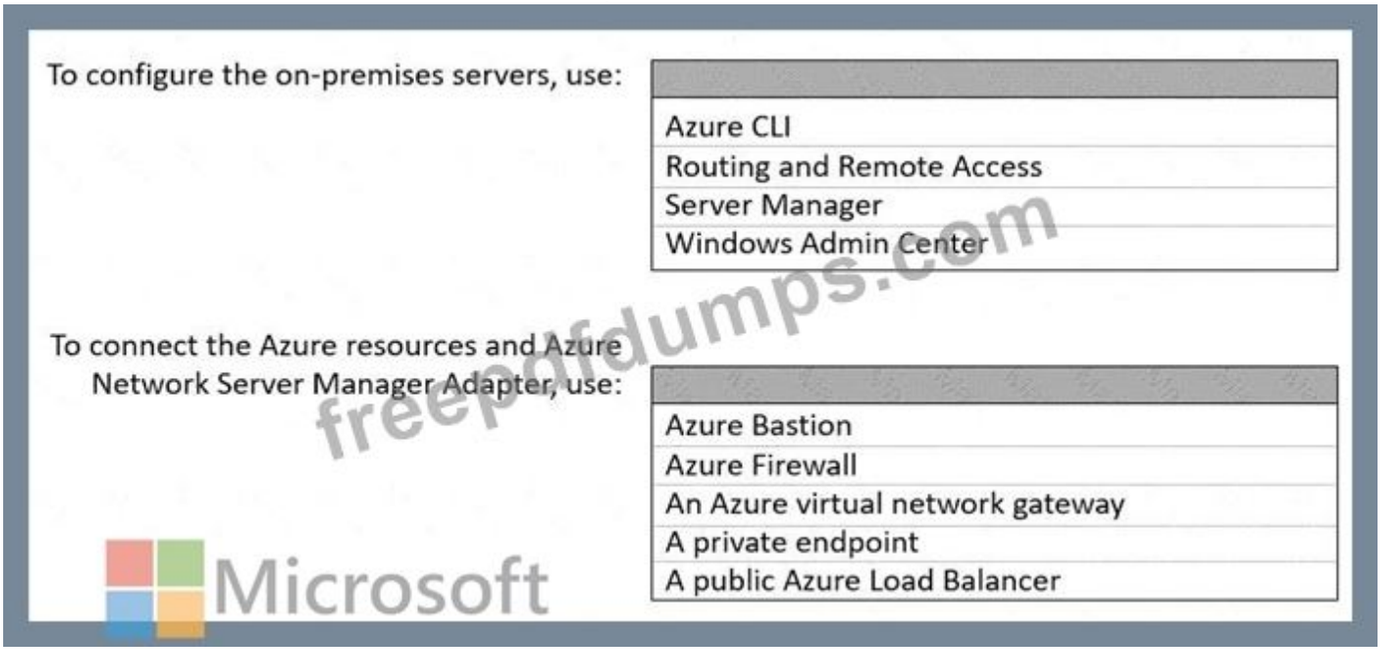

Answer:

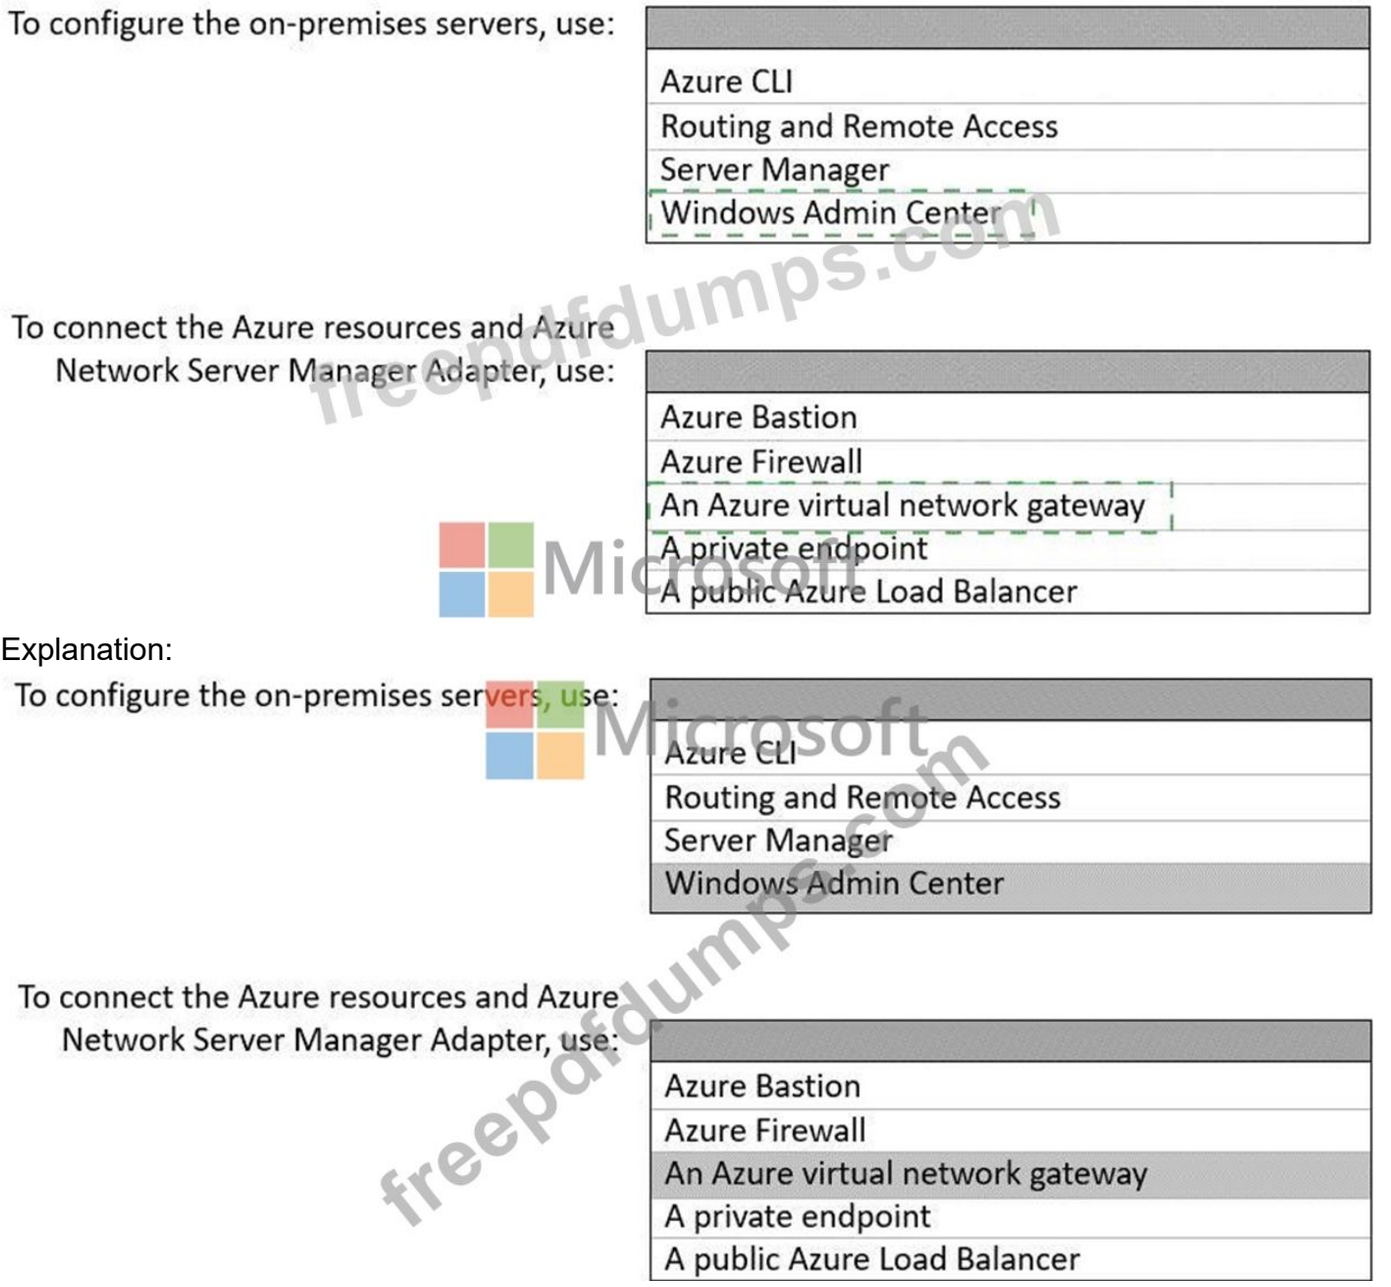

Explanation:

In the Administering Windows Server Hybrid Core Infrastructure materials, Azure Network Adapter (ANA) is presented as a Windows Admin Center (WAC) capability used to connect on-premises Windows Server computers to Azure. The guidance states that Windows Admin Center provides an Azure Network Adapter workflow that creates a point-to-site (P2S) VPN from the selected on-premises server to an Azure virtual network. The server-side configuration (VPN profile, certificates, and routes) is pushed directly from WAC; therefore, you use Windows Admin Center-not Server Manager, Azure CLI, or RRAS-to configure the on-premises servers.

On the Azure side, the same module explains that the P2S connection terminates on an Azure VPN Gateway deployed in the target virtual network. ANA either uses an existing gateway or helps you provision one by ensuring a GatewaySubnet exists and configuring Point-to-Site settings. The study content emphasizes that "Azure Network Adapter establishes a point-to-site VPN to an Azure VNet by using an Azure VPN gateway; other Azure services such as Bastion,

Firewall, Private Endpoints, or Load Balancers are not used for this tunnel." Consequently, the required Azure component is an Azure virtual network gateway.
Putting it together: to set up ANA you configure the on-premises servers with Windows Admin Center, and to connect to Azure you use an Azure virtual network gateway as the VPN termination point.

NEW QUESTION: 82
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.

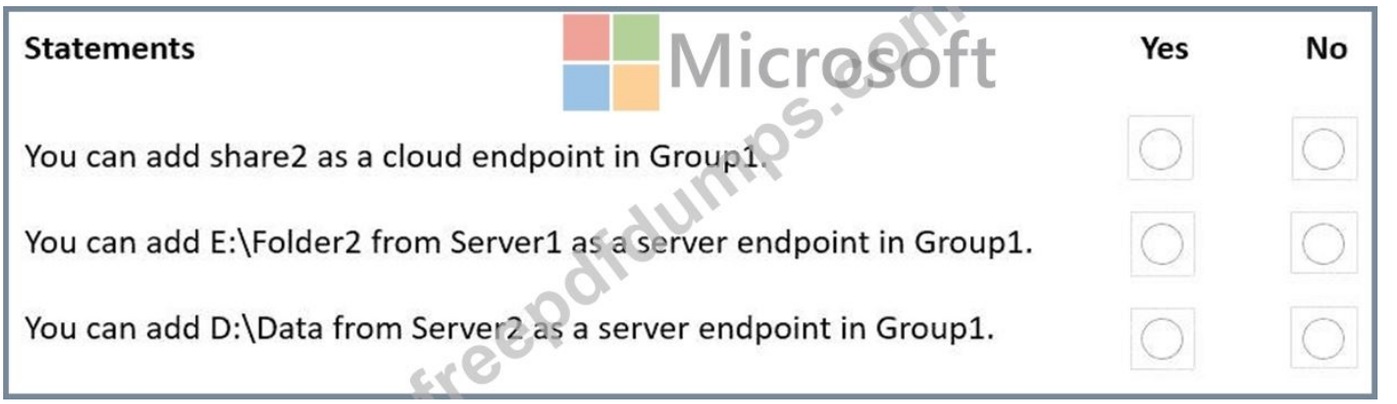

Answer:

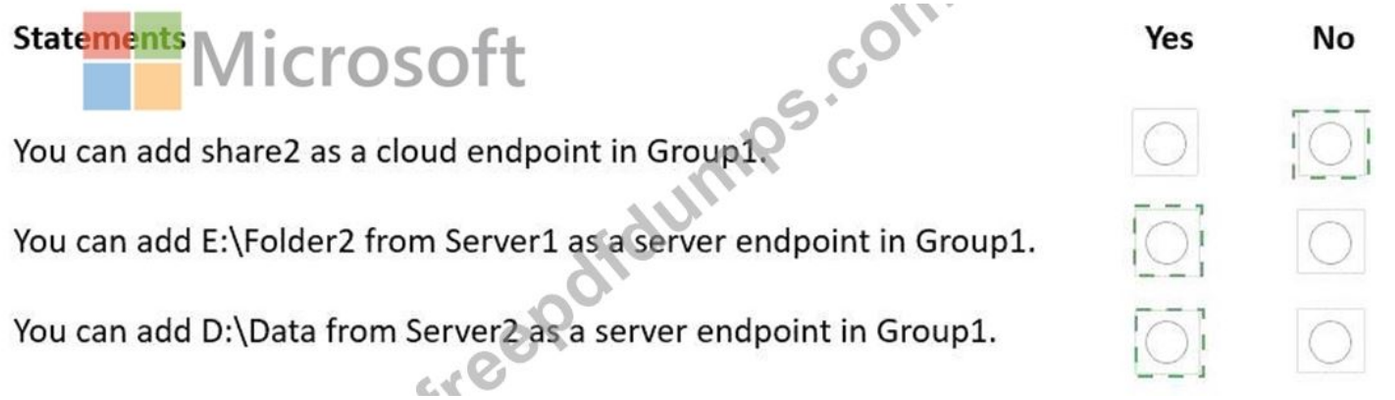

Explanation:

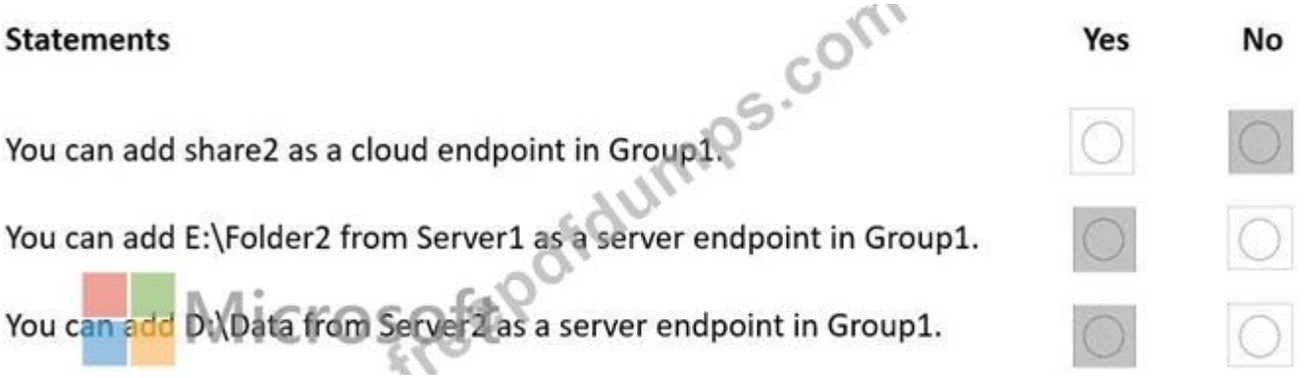

The Azure File Sync design guidance in Administering Windows Server Hybrid Core Infrastructure explains that a sync group defines one sync topology and is anchored by exactly one cloud endpoint (an Azure file share). The material states: "Each sync group contains a single cloud endpoint... Additional Azure file shares must be placed in separate sync groups." Consequently, after Group1 uses share1 as its cloud endpoint, share2 cannot be added to the same sync group, which makes the first statement No.

For server endpoints, the guide notes: "A server endpoint represents a specific path on a registered Windows Server... A server can host multiple server endpoints, including in the same sync group, provided paths don't overlap (no parent/child or identical paths)." Because D:\Folder1 is already in Group1, adding E:\Folder2 from the same Server1 is supported and non-overlapping, so the second statement is Yes.

Finally, multi-server synchronization is a core capability: "Multiple servers can be added as server endpoints to the same sync group to enable branch-to-branch and server-to-cloud sync; contents

are merged in the Azure file share namespace." Therefore, adding D:\Data from Server2 as another server endpoint in Group1 is fully supported, making the third statement Yes.
These rules satisfy the scenario: one cloud endpoint per sync group, multiple non-overlapping server endpoints (including multiple from the same server), and multi-server participation in the same group.

NEW QUESTION: 83
You have an Azure subscription that contains the virtual machines shown in the following table.

Name	Operating system
VM1	Windows Server 2022 Datacenter: Azure Edition
VM2	Windows Server 2022 Datacenter: Azure Edition Core
VM3	Windows Server 2022 Datacenter
VM4	Windows Server 2019 Datacenter

You plan to implement Azure Automanage for Windows Server.
You need to identify the operating system prerequisites.
Which virtual machines support Hotpatch, and which virtual machines support SMB over QUIC?
To answer select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

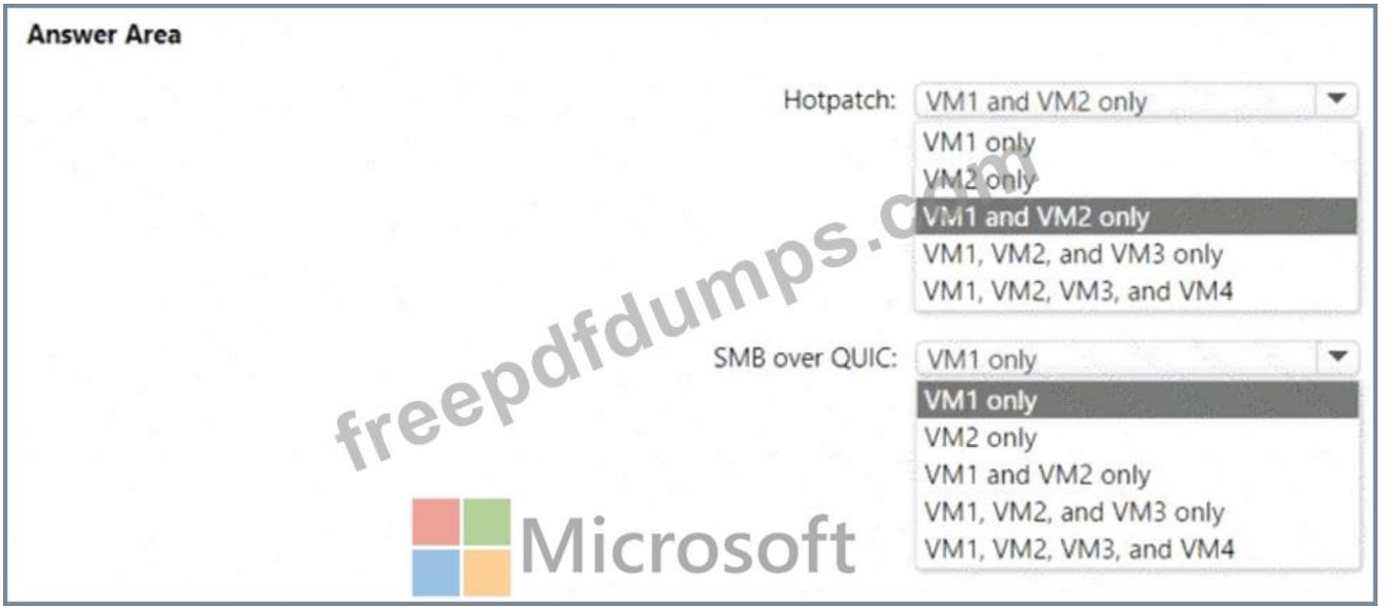

Answer:

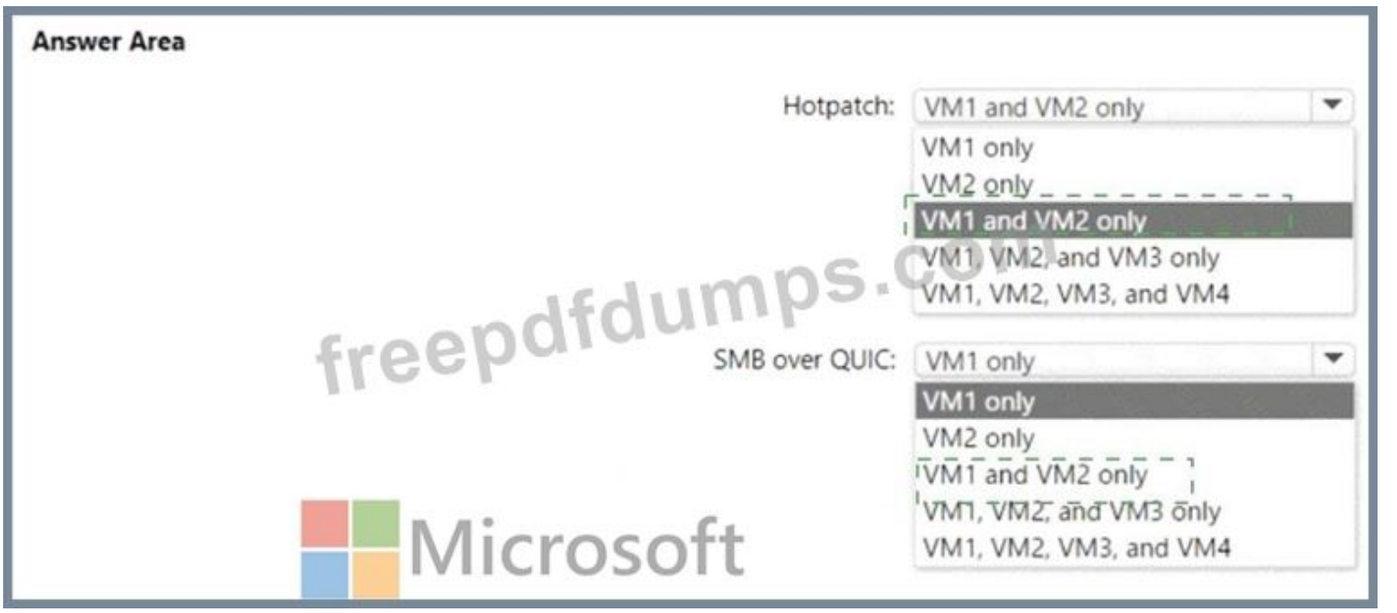

Explanation:
Hotpatch: VM1 and VM2 only
SMB over QUIC: VM1 and VM2 only
The Administering Windows Server Hybrid Core Infrastructure guidance for Azure Automanage and Windows Server 2022 features states that Hotpatch is available only for Azure VMs that run Windows Server 2022 Datacenter: Azure Edition. The materials describe Hotpatch as an Azure capability that
"applies security updates to Windows Server 2022 Azure Edition without requiring a reboot in most cases," and clarify that it is supported on both Desktop Experience and Server Core images of Azure Edition used in Azure IaaS. Therefore, only VM1 (Windows Server 2022 Datacenter: Azure Edition) and VM2 (Windows Server 2022 Datacenter: Azure Edition Core) meet the Hotpatch prerequisite. Regular Windows Server 2022 Datacenter (VM3) and Windows Server 2019 Datacenter (VM4) are not eligible for Hotpatch.
For file services, the same exam content explains that SMB over QUIC is a Windows Server 2022 feature limited to Azure Edition when acting as the SMB server endpoint: it "provides SMB encryption and TLS 1.3 over UDP (QUIC) to enable secure, VPN-less access to SMB shares" and is available on Windows Server
2022 Datacenter: Azure Edition. Consequently, only VM1 and VM2 can host SMB over QUIC. Standard Windows Server 2022 Datacenter (VM3) and Windows Server 2019 Datacenter (VM4) do not provide the SMB over QUIC server capability.
Thus, the correct selections are VM1 and VM2 only for both Hotpatch and SMB over QUIC.

NEW QUESTION: 84
You need to meet the technical requirements for User1. The solution must use the principle of least privilege.
What should you do?
A. Add Users1 to the Server Operators group in contoso.com.
B. Create a delegation on contoso.com.

C. Add Users1 to the Account Operators group in contoso.com.
D. Create a delegation on OU3.
Answer: D (LEAVE A REPLY)
n the Administering Windows Server Hybrid Core Infrastructure guidance for managing AD DS, Microsoft emphasizes using OU-level delegation to satisfy administrative needs while adhering to the principle of least privilege. The documentation explains that the Delegate Control wizard on an OU lets you grant a user or group only the specific permissions required for common tasks, including "Modify the membership of a group". This grants the write permission to the member attribute on group objects contained in that OU, without giving broader account-management rights across the domain.
By contrast, placing a user in Account Operators or Server Operators provides elevated, domain-wide capabilities far beyond what is required. Account Operators can create, delete, and modify many account types across the domain (except for protected admin accounts), which violates least-privilege for a task that only needs to change group membership in one OU. Server Operators is unrelated to group membership and is intended for server administration tasks. Creating a delegation at the domain root would similarly be excessive because it applies broadly to all containers and OUs.
Therefore, to meet the requirement "Ensure that User1 can manage the membership of all the groups in Contoso\OU3," you should delegate control on OU3 and assign the built-in task "Modify the membership of a group" to User1, achieving the minimal permissions necessary.

NEW QUESTION: 85
Task 3
You need to create 3 user named Admin1 in contoso.com. Admin1 must be able to back up and restore files on SRV1. The solution must use principle of the least privilege.
Answer:
See the solution of this Task below.
Explanation:
TASK 3
Objective:
Create a user named Admin1 in contoso.com.
Admin1 must be able to back up and restore files on SRV1.
Follow the principle of least privilege.
Step-by-Step Guide
Step 1: Create the User Account
Log in to a Domain Controller (e.g., DC1) with appropriate admin rights.
Open Active Directory Users and Computers (dsa.msc).
In the contoso.com domain:
Right-click the Users container or another OU where you want to create the account.
Select New > User.
Enter the following:

First name: Admin1
User logon name: Admin1
Click Next and set a password (ensure it meets the domain's password policy).
Configure password options (e.g., User must change password at next logon, if required).
Click Finish.
Step 2: Grant Backup and Restore Rights on SRV1
By default, Backup Operators have the ability to back up and restore files (without giving full admin rights).
Log in to SRV1 (the target server).
Open Computer Management (compmgmt.msc).
In the left pane, expand:
System Tools > Local Users and Groups > Groups.
Find and double-click the Backup Operators group.
Click Add.
In the Select Users, Computers, Service Accounts, or Groups window:
Type Admin1.
Click Check Names to resolve the user.
Click OK to add Admin1 to the group.
Click OK again to close the Backup Operators group properties.
Step 3: Verify Access
Log in as Admin1 on SRV1 and test performing backup and restore operations using tools like Windows Server Backup.
Since Backup Operators can back up and restore data but do not have full administrative privileges, this follows the least privilege principle.
Additional Notes
If you prefer using PowerShell, you can add the user to the group like this on SRV1:
Add-LocalGroupMember -Group "Backup Operators" -Member "contoso\Admin1"

NEW QUESTION: 86
Your network contains an Active Directory Domain Services (AD DS) domain.
You have a Group Policy Object (GPO) named GPO1 that contains Group Policy preferences.
You plan to link GPO1 to the domain.
You need to ensure that the preference in GPO1 apply only to domain member servers and NOT to domain controllers or client computers. All the other Group Policy settings in GPO1 must apply to all the computers.
The solution must minimize administrative effort.
Which type of item level targeting should you use?
A. Domain
B. Operating System
C. Security Group
D. Environment Variable

Answer: B (LEAVE A REPLY)
Reference:
In the Windows Server hybrid administration curriculum, Group Policy Preferences (GPP) support Item-level Targeting (ILT), which lets you apply a preference only when specific conditions are met, while the remainder of the GPO (regular policy settings) continues to apply wherever the GPO is linked. The guide explains that ILT can evaluate many attributes, including Operating System, where you can match by version, architecture, and crucially Product type. The product type values distinguish Workstation, Server, and Domain Controller. Using the Operating System targeting item to set Product type = Server ensures the preference applies to member servers only, because domain controllers report a different product type (Domain Controller) and client computers report Workstation.
This approach satisfies the requirement to link GPO1 at the domain level so that "all the other Group Policy settings in GPO1 apply to all computers," while constraining only the GPP items to servers. It also meets the
"minimize administrative effort" goal because it requires no OU restructuring and no additional security groups or WMI filters. In short, configure ILT on the specific GPP items within GPO1 using Operating System # Product type: Server; DCs and clients will not meet the targeting condition, so the preferences won't apply to them, but the rest of the GPO will.

NEW QUESTION: 87
Your network contains an Active Directory Domains Services (AD DS) domain named contoso.com. You implement a central store.
You create a new Group Policy Object (GPO) named GP01.
When you attempt to edit GP01, you see the settings shown in the exhibit. (Click the Exhibit tab.)
You need to ensure that all settings are available.
Solution: You modify the properties of GPO1.
Does this meet the goal?
A. No
B. Yes
Answer: (SHOW ANSWER)

NEW QUESTION: 88
Task 10
You need to ensure that SRV1 only leases IP addresses from the range of 192.168.1.190 to 192.168.1.200 to computers that have a MAC address that starts with aabb.
Answer:
See the solution of this Task below.
Explanation:
Objective:
Configure the DHCP server SRV1 to lease IP addresses only to computers with MAC addresses starting with AABB in a specific range.

Step-by-Step Guide

Step 1: Open DHCP Management Console

* Log in to SRV1 with Domain Admin or DHCP Admin privileges.
* Open DHCP Manager:
* Press Windows + R, type dhcpmgmt.msc, and press Enter.

Step 2: Create a New DHCP Scope

* In the DHCP console, expand SRV1.
* Right-click IPv4 and select New Scope.
* The New Scope Wizard opens.

Step 3: Configure the Scope

* Name:
* Enter a name (e.g., MAC-Filtered Scope).
* Click Next.
* IP Address Range:
* Start IP: 192.168.1.190
* End IP: 192.168.1.200
* Subnet mask: as appropriate (e.g., 255.255.255.0).
* Click Next.
* Add Exclusions:
* None needed unless you want to reserve certain addresses.
* Click Next.
* Lease Duration:
* Set as needed, default is usually fine.
* Click Next.
* Configure DHCP Options:
* You can skip or configure as needed (gateway, DNS, etc.).
* Click Next.
* Activate Scope:
* Click Yes to activate it.

Step 4: Configure MAC Address Filtering (Allow List)

* In the DHCP console, expand the scope you created.
* Right-click Filters under the scope and choose New Filter.
* Enter the MAC address pattern to match devices with MAC addresses starting with AABB:
* MAC Address: AABB*
* Description: e.g., Allow devices starting with AABB.
* Click Add.

Step 5: Enable Allow Filters

* Right-click Filters under the scope and select Enable.
* Ensure that only devices matching the AABB pattern will receive leases.

Step 6: Test and Verify

* Use a test client with a MAC address starting with AABB to ensure it receives an IP address in the
192.168.1.190-192.168.1.200 range.
* Use ipconfig /renew on the client, or check the DHCP leases in the Address Leases section.

NEW QUESTION: 89
You have a Windows Server container host named Server 1 and a container image named Image1.
You need to start a container from image1. The solution must run the container on a Hyper-V virtual machine.
Which parameter should you specify when you run the docker run command?
A. --expose
B. --privileged
C. --runtime
D. --entrypoint
E. --isolation
Answer: E (LEAVE A REPLY)
Reference:
In Windows Server container scenarios, process isolation shares the host kernel, while Hyper-V isolation runs each container inside a lightweight Hyper-V VM that provides a stronger boundary. The Windows Server Hybrid Core Infrastructure guidance states that when you must run a Windows container with a VM boundary, you start it using Hyper-V isolation. Operationally, this is done at run time with the Docker CLI by specifying the isolation mode: docker run --isolation=hyperv .... Other parameters don't meet the requirement: --expose only publishes ports, --privileged is a Linux concept not used for Windows security boundaries, --runtime selects the OCI runtime (relevant to Linux), and --entrypoint overrides the default process but does nothing for isolation. Using --isolation=hyperv ensures the container launches on a minimal Hyper-V partition created by the Windows container host, satisfying scenarios that require strict separation, kernel mismatch tolerance, or enhanced defense-in-depth-exactly what the requirement "run the container on a Hyper-V virtual machine" describes. This aligns with the exam's emphasis on selecting the proper Windows container isolation mode to meet security and compatibility goals during deployment and operations.

NEW QUESTION: 90
You have an on-premises network that is connected to an Azure virtual network by using a Site-to-Site VPN.
Each network contains a subnet that has the same IP address space. The on-premises subnet contains a virtual machine.
You plan to migrate the virtual machine to the Azure subnet.
You need to migrate the on premises virtual machine to Azure without modifying the IP address. The solution must minim administrative effort.

What should you implement before you perform the migration?

A. Azure Extended Network

B. Azure Virtual Network NAT

C. Azure Application Gateway

D. Azure virtual network peering

Answer: A (LEAVE A REPLY)

The AZ-800 materials discuss options for migrating on-premises workloads to Azure when IP address preservation is required and subnet overlaps exist across a Site-to-Site VPN. The guidance explains that Azure Extended Network allows you to "stretch an on-premises subnet into Azure so a VM can retain its original IP address after migration" and that it is specifically intended to "avoid renumbering during lift-and-shift scenarios where overlapping address space or application dependencies prevent changing IPs." In contrast, services like VNet peering or Application Gateway do not solve overlapping address space or IP preservation for a VM, and NAT in a VNet translates traffic but does not allow the VM to keep its original IP address. The hybrid curriculum highlights that Extended Network "minimizes administrative effort by maintaining existing addressing and DNS records during migration," which matches the requirement to migrate the on-premises VM to Azure without modifying the IP address and with minimal changes to the environment. Therefore, implementing Azure Extended Network before the migration is the correct and least-effort solution.

NEW QUESTION: 91

Your network contains an Active Directory Domain Services (AD DS) domain. The domain contains a user named User1 and the servers shown in the following table.

Name	Server role	DHCP scope
Server1	DHCP Server	Scope1, Scope2
Server2	DHCP Server	Scope3, Scope4

You need to ensure that User1 can manage only Scope1 and Scope3. What should you do?

A. Add User1 to the DHCP Administrators group on Server1 and Server2.

B. Implement IP Address Management (IPAM).

C. Add User1 to the DHCP Administrators domain local group.

D. Implement Windows Admin Center and add connections to Server1 and Server2.

Answer: (SHOW ANSWER)

The Windows Server hybrid core curriculum describes IPAM as follows: "IP Address Management (IPAM) provides centralized administration and role-based access control (RBAC) for DHCP and DNS." It further specifies: "IPAM enables delegation such as DHCP Scope Administrator, allowing an operator to manage specific scopes on selected DHCP servers without granting full server administration." By contrast, the built- in groups provide coarse permissions: "Adding a user to DHCP Administrators (local or domain) grants broad administrative rights over the DHCP server and all scopes." Windows Admin Center does not introduce per-scope RBAC for DHCP; it simply surfaces the existing server tools. In the scenario, the requirement is that

User1 can manage only Scope1 and Scope3-that is scope-level delegation across two servers. The only option among the choices that provides granular, scope-scoped rights is IPAM, where you assign the DHCP Scope Administrator role to User1 for Scope1 on Server1 and Scope3 on Server2, leaving other scopes unmanaged by that user.

NEW QUESTION: 92
You need to ensure that VM3 meets the technical requirements.
What should you install first?
A. Enhanced Storage
B. File Server Resource Manager (FSRM)
C. Windows Standards-Based Storage Management
D. the iSNS Server service
Answer: B (LEAVE A REPLY)
To meet the requirement that VM3 be configured for per-folder quotas, the Windows Server file services module directs you to install File Server Resource Manager (FSRM). The course content explains: "FSRM provides quota management, file screening, and storage reporting. Quotas can be applied to volumes and to specific folders, enabling administrators to control the space consumed." None of the other features listed deliver folder-level quota capability: Enhanced Storage and Windows Standards-Based Storage Management relate to storage management/SMI-S, and iSNS Server concerns iSCSI discovery services. Therefore, the first step is to install FSRM on VM3; after installation, you can create per-folder quota templates and assignments to enforce the desired limits.

NEW QUESTION: 93
You have an on-premises Active Directory Domain Services (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 membership.
D. in the Azure AD tenant create a Microsoft 365 group
Answer: (SHOW ANSWER)
The hybrid identity guidance for Windows Server and Azure emphasizes using Azure AD (Microsoft Entra ID) security groups for unified authorization across Azure Files (Azure Storage with Azure AD-based access control) and on-premises SMB shares-when Group writeback is enabled in Azure AD Connect. The documented pattern is: create the group in Azure AD with assigned membership so you can grant Azure roles
/share permissions to storage2\share2 using Azure AD identities. Because Group writeback is enabled, that Azure AD security group is written back to on-premises AD DS as a universal security group, allowing you to use the same group on NTFS/share permissions for Server1\Share1. Members can be synchronized AD DS users; adding them in Azure AD (assigned membership) ensures predictable inclusion for both resources.
Creating the group only in AD DS (Option A) can sync to Azure, but the hybrid guidance highlights Azure- originated groups with writeback when you want one source to drive both cloud and on-prem authorization and to ensure compatibility with Azure Files' Azure AD-based access checks. Dynamic membership (Option C) is unnecessary here and may not align with file access scenarios. Microsoft 365 groups (Option D) are not intended for file share ACLs/RBAC in this context. Thus, the verified solution is to create an Azure AD security group with assigned membership, taking advantage of group writeback for on-premises use.

NEW QUESTION: 94
Your on-premises network contains a single-domain Active Directory Domain Services (AD DS) forest. You have an Azure AD tenant named contoso.com. The AD DS forest syncs with the Azure AD tenant by using Azure AD Connect.
You need to ensure that users in the forest that have a custom attribute of NoSync are excluded from synchronization.
How should you configure the Azure AD Connect cloudFiltered attribute, and which tool should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

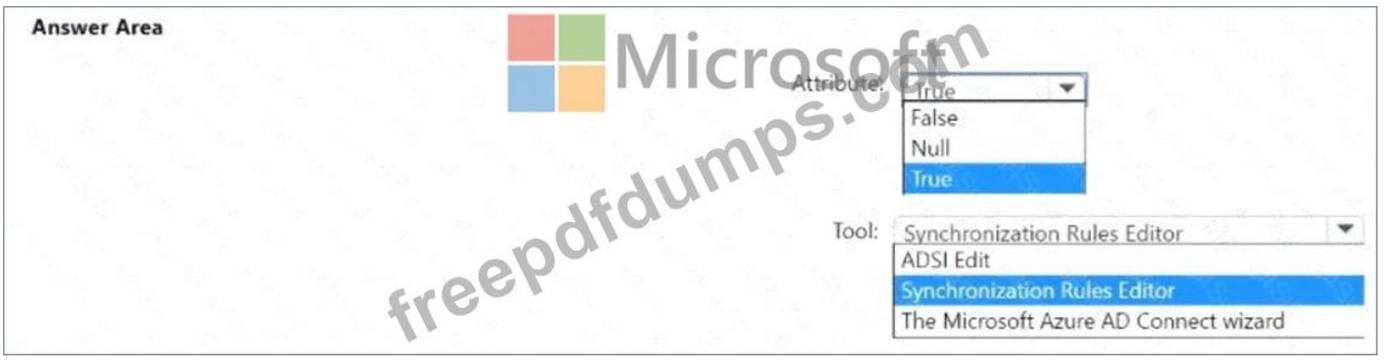

Answer:

Explanation:

In Azure AD Connect-based hybrid identity, object filtering can be enforced with the built-in cloudFiltered metaverse flag. The Administering Windows Server Hybrid Core Infrastructure materials explain that

"cloudFiltered is a Boolean in the metaverse that determines whether an object is eligible for export to Azure AD; when set to True, the connector space object is excluded from synchronization to Azure AD." The recommended approach for attribute-based filtering is to create a custom inbound synchronization rule that scopes the users you want to exclude and flows a constant value of True into cloudFiltered. The guide states: "Use the Synchronization Rules Editor to define an inbound rule from Active Directory to the metaverse. Add a scoping filter (for example, a custom attribute equals 'NoSync') and configure an attribute flow mapping cloudFiltered # Constant(True). Objects matching the scope will not be exported." Because you need to exclude any user whose custom attribute equals NoSync, you create an inbound rule targeting user objects, add the scoping filter on that custom attribute, and then set cloudFiltered = True (constant). This is done with the Synchronization Rules Editor, not ADSI Edit or the Azure AD Connect wizard, since the wizard doesn't author custom attribute flows and ADSI Edit doesn't manage sync rules.

NEW QUESTION: 95

Task 5

You have an application that is copied to a folder named C:\app on SRV1. C:\app also contains also a Dockerfile for the app.
On SRV1. you need to create a container image for the application by using the Dockerfile. The container image mutt be named app1.

Answer:

See the solution of this Task below.

Explanation:

Explore

To create a container image named app1 for your application using the Dockerfile in the C:\app directory on SRV1, follow these steps:

Step 1: Open PowerShell or Command Prompt First, open PowerShell or Command Prompt on SRV1.

Step 2: Navigate to the Application Directory Change to the directory where your application and Dockerfile are located:

cd C:\app

Step 3: Build the Container Image Use the docker build command to create the container image. The -t flag tags the image with the name app1:

docker build -t app1 .

The period . at the end of the command tells Docker to use the Dockerfile in the current directory.

Step 4: Verify the Image Creation After the build process completes, verify that the image app1 has been created successfully by listing all images:

docker images

You should see app1 in the list of images.

Step 5: Use the Image Now, you can use the image app1 to run containers or push it to a container registry if needed.

By following these steps, you'll have created a Docker container image named app1 using the Dockerfile located in C:\app on SRV11. Ensure that Docker is installed on SRV1 and that you have the necessary permissions to execute these commands.

NEW QUESTION: 96

You need to meet the technical requirements for Server1. Which users can currently perform the required tasks?

A. Admin1 only

B. Admin3 only

C. Admin1 and Admin3 only

D. Admin1 Admin2. and Admm3

Answer: (SHOW ANSWER)

In the AZ-800 "Administering Windows Server Hybrid Core Infrastructure" objectives for Active Directory, server promotion is governed by forest/domain administrative roles. The materials state that promoting a member server to a domain controller in a given domain requires membership in either the Enterprise Admins group or the Domain Admins group of the target domain. The Configuration and Domain naming contexts that DCPromo touches (NTDS settings, SYSVOL/DFS-R readiness, DC computer account, and associated service SPNs) are protected so that "Enterprise Admins have full rights forest-wide, and Domain Admins have full rights within their respective domain." In this case, the requirement is to promote Server1 to a domain controller in canada.contoso.com. From the identities table:

* Contoso\Admin1 is a member of Enterprise Admins (forest-wide authority).
* Canada\Admin3 is a member of Canada\Domain Admins (authority within canada.contoso.com).
* Contoso\Admin2 is Domain Admins (contoso.com) only, which does not grant administrative authority in the canada.contoso.com child domain.

Therefore, the users who can currently perform the required task for Server1 are Admin1 and Admin3.

NEW QUESTION: 97

Task 1

You need to create a group-managed service account (gMSA) named gMSA1 and make gMSA1 available on SRV1.

Answer:

See the solution of this Task below.

Explanation:

To create a group-managed service account (gMSA) named gMSA1 and make it available on SRV1, you can follow these steps:

Step 1: Create the Key Distribution Services Root Key First, you need to create the KDS Root Key, which is required for the gMSA to function. You can do this with the following PowerShell command:

Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))

Note: The -EffectiveTime parameter is set to 10 hours in the past to ensure immediate effect.

Step 2: Create the gMSA Next, use the New-ADServiceAccount cmdlet to create the gMSA:

New-ADServiceAccount -Name gMSA1 -DNSHostName gmsa1.domain.com -PrincipalsAllowedToRetrieveManagedPassword SRV1$

Replace domain.com with your actual domain name.

Step 3: Install the gMSA on SRV1 Now, you need to install the gMSA on the server SRV1. Run the following command on SRV1:
Install-ADServiceAccount -Identity gMSA1
Step 4: Test the gMSA To ensure that the gMSA is installed correctly and ready for use, perform a test using:
Test-ADServiceAccount -Identity gMSA1
If the test returns True, the gMSA is correctly installed and ready for use on SRV1.
Step 5: Configure the Service to Use the gMSA Finally, configure the service that requires the gMSA to use gMSA1 by setting the service's logon account to domain\gMSA1$ and leave the password field blank.
This will create and make the gMSA gMSA1 available on SRV1. Ensure that you have the necessary permissions and that SRV1 is properly joined to the domain before proceeding with these steps123.

NEW QUESTION: 98
Task 7
SRV1 contains a virtual machine named VM1.
You need attach c:\vhds\DiskI.vhdxtoVMI.The solution must ensure that Disk1 can be expanded dynamically when VM1 runs.
Answer:
See the solution of this Task below.
Explanation:
TASK 7
Objective:
Attach the VHDX file c:\vhds\Disk1.vhdx to VM1 so it can dynamically expand when VM1 is running.
Step-by-Step Guide
Step 1: Verify the VHDX File Type
Dynamic expansion means the virtual disk type should be dynamic (not fixed).
Let's verify the disk type of Disk1.vhdx:
powershell
Copy
Get-VHD -Path "c:\vhds\Disk1.vhdx"
In the output, ensure that VhdType shows Dynamic.
If it's not dynamic, convert it:
powershell
Copy
Convert-VHD -Path "c:\vhds\Disk1.vhdx" -DestinationPath "c:\vhds\Disk1_dynamic.vhdx" -VHDType Dynamic
Step 2: Attach the VHDX to VM1
Use PowerShell to add the disk to VM1:

powershell
Copy
Add-VMHardDiskDrive -VMName "VM1" -Path "c:\vhds\Disk1.vhdx"
Or do it via Hyper-V Manager:
Open Hyper-V Manager.
Select VM1 and go to Settings.
In the left pane, select SCSI Controller (recommended for hot-add).
Click Add Hard Drive.
Browse and select c:\vhds\Disk1.vhdx.
Click Apply and OK.
Step 3: Verify Hot-Add Support (Optional)
If VM1 is running Windows Server 2012 or later and uses a SCSI Controller, the VHDX can be added without shutting down the VM.
Step 4: Verify the Disk in the VM
Inside VM1, open Disk Management (diskmgmt.msc).
The newly attached disk should appear as unallocated.
Initialize it, create a volume, and format if needed.
Summary
VHDX file type: Must be dynamic for expanding.
Hot-adding supported if attached via SCSI Controller and OS supports it.
Now Disk1.vhdx is attached to VM1 and can expand dynamically.

NEW QUESTION: 99
Your network contains an Active Directory Domain Services (AD DS) domain. The domain contains the domain controllers shown in the following table.

Name	Description
DC1	PDC emulator, RID master, and global catalog server
DC2	Infrastructure master and domain naming master
DC3	Schema master
RODC1	Read-only domain controller (RODC)

You need to ensure that if an attacker compromises the computer account of RODC1, the attacker cannot view the Employee-Number AD DS attribute. Which partition should you modify?
A. configuration
B. global catalog
C. domain
D. schema
Answer: D (LEAVE A REPLY)
The domain services security section describes the RODC Filtered Attribute Set (FAS), which prevents selected attributes from replicating to Read-Only Domain Controllers. The curriculum

explains: "Attributes that must not be exposed on RODCs are added to the RODC filtered attribute set. This setting is defined in the schema partition; once an attribute is flagged as CONFIDENTIAL and added to FAS, it is not replicated to any RODC." Because the attacker may have compromised RODC1's computer account, protection requires ensuring sensitive attributes (e.g., employeeNumber) are never present on that RODC. Changes for FAS membership and the confidential bit are schema-level operations and apply forest-wide. Modifying the domain or configuration partitions will not change which attributes replicate to RODCs, and the global catalog is not a writable partition; it aggregates a subset of attributes but honors the schema's FAS configuration. Therefore, to ensure Employee-Number cannot be viewed if RODC1 is compromised, you modify the schema to place the attribute in the RODC filtered attribute set (and, if required, mark it confidential).

NEW QUESTION: 100
You need to configure remote administration to meet the security requirements. What should you use?
A. just in time (JIT) VM access
B. Azure AD Privileged Identity Management (PIM)
C. the Remote Desktop extension for Azure Cloud Services
D. an Azure Bastion host
Answer: A (LEAVE A REPLY)
In the Administering Windows Server Hybrid Core Infrastructure materials (hybrid security and IaaS management), Just-In-Time (JIT) VM access from Microsoft Defender for Cloud is the prescribed way to require approval-based, time-bound Remote Desktop access to Azure VMs. The guide explains that JIT
"locks down inbound traffic to management ports (for example, TCP/3389) and opens them only on request, for a limited time and only from approved source IPs." Administrators request access; upon approval, Defender for Cloud creates a temporary NSG rule that expires automatically-you specify the maximum allowed window (e.g., 2 hours) and the ports. This matches the requirement: "Ensure that server administrators request approval before they can establish a Remote Desktop connection to an Azure virtual machine. If the request is approved, the connection must be established within two hours." Alternatives don't meet this: PIM governs Azure roles, not VM RDP port exposure; Azure Bastion provides secure RDP/SSH over TLS without public IPs but doesn't provide approval/time-boxed gating; the Remote Desktop extension is for classic Cloud Services and not for policy-driven approval windows. JIT is the least-privilege, policy- enforced solution aligned with the exam's hybrid security objectives.

NEW QUESTION: 101
Task 10
You use a Group Policy preference to map \\dd.contoso.com\instal1 as drive H for all users. If a user already has an existing drive mapping for H. the new drive mapping must take precedence.
Answer:

See the solution of this Task below.
Explanation:
To map \\dd.contoso.com\instal1 as drive H for all users using Group Policy Preferences and ensure that the new drive mapping takes precedence over any existing mappings, follow these steps:
Step 1: Open Group Policy Management Console Open the Group Policy Management Console (GPMC) on a machine that has administrative privileges over the domain.
Step 2: Create or Edit a GPO Create a new Group Policy Object (GPO) or edit an existing one that applies to the users who need the drive mapping.
Step 3: Navigate to Drive Mappings In the GPO Editor, navigate to:
User Configuration -> Preferences -> Windows Settings -> Drive Maps
Step 4: New Drive Mapping Right-click on Drive Maps and select New -> Mapped Drive.
Step 5: Configure Drive Mapping In the New Drive Properties window, configure the following settings:
* Action: Select Replace. This action will overwrite any existing mappings with the same drive letter.
* Location: Enter the UNC path \\dd.contoso.com\instal1.
* Drive Letter: Choose H: from the drop-down menu.
* Reconnect: Check this option if you want the drive mapping to persist across logon sessions.
* Label As: Optionally, provide a label for the drive mapping.
* Hide/Show this drive: Set according to your preference.
* Hide/Show all drives: Set according to your preference.
Step 6: Common Tab Go to the Common tab and configure the following:
* Run in logged-on user's security context (user policy option): Check this option.
* Item-level targeting: Click on Targeting and set up any specific criteria if needed.
Step 7: Apply the GPO Click Apply and then OK to save the drive mapping configuration.
Step 8: Link the GPO Link the GPO to an Organizational Unit (OU) or domain that contains the users who should receive the drive mapping.
Step 9: Update Group Policy Instruct users to log off and log back on, or use the gpupdate /force command to refresh Group Policy on their computers.

NEW QUESTION: 102
Task 1
You need to ensure that DC2 is the schema master for contoso.com.
Answer:
See the solution of this Task below.
Explanation:
Step-by-Step Guide: Seizing/Transferring the Schema Master Role to DC2
Step 1: Log in to DC2
Use an account that is a member of the Schema Admins, Enterprise Admins, and Domain Admins groups.

Step 2: Register the Schema Snap-in
The Schema snap-in is not loaded by default.
Open Command Prompt as Administrator.
Type the following command to register the schema management DLL:
powershell
Copy
regsvr32 schmmgmt.dll
Step 3: Open MMC (Microsoft Management Console)
Press Windows + R, type mmc, and hit Enter.
In MMC, go to File > Add/Remove Snap-in.
Select Active Directory Schema, then click Add > OK.
Step 4: Connect to DC2
In the Active Directory Schema console, right-click Active Directory Schema and select Change Active Directory Domain Controller.
In the dialog box, select DC2 and click OK.
This will connect the console to DC2.
Step 5: Transfer the Schema Master Role
Right-click Active Directory Schema again and select Operations Master.
In the Change Schema Master dialog box, confirm that DC2 is shown as the target.
Click the Change button to transfer the Schema Master role to DC2.
Click Yes when prompted to confirm the transfer.
Step 6: Verify the Transfer
In the same dialog box, ensure that DC2 is now listed as the Schema Master.
Optionally, run the following command in PowerShell to verify:
netdom query fsmo
The Schema Master should now be DC2.

NEW QUESTION: 103
Your network contains an on-premises Active Directory Domain Services (AD DS) domain. The domain contains a user named User1 and the servers shown in the following table.

Name	Operating system
Server1	Windows Server 2016
Server2	Windows Server 2022
Backup1	Windows Server 2019

User1 is a member of the Protected Users security group.
User1 performs the following actions:
* From Server1, establishes a remote PowerShell session on Server2
* From the PowerShell session on Server2, attempts to access a resource on Backup1 The request to access the resource on 8ackup1 is denied.

You need to ensure that User1 can access the resources on Backup1 by using the PowerShell session on Server2. The solution must follow the principle of least privilege and minimize administrative effort.
What should you configure?
A. Kerberos delegation (unconstrained)
B. CredSSP
C. PSSessionConfigu ration by using RunAs
D. resource-based Kerberos constrained delegation
Answer: D (LEAVE A REPLY)
In Windows remoting, hopping from one server to another (Server1 # Server2 # Backup1) triggers the
"second-hop" problem because the user's TGT is not forwarded. The AZ-800 material explains that members of Protected Users cannot use NTLM, DES/RC4, or unconstrained delegation, and their credentials cannot be cached or delegated. Therefore, CredSSP or unconstrained delegation cannot be used. The guide prescribes Kerberos constrained delegation to allow a middle-tier server to act on the user's behalf, and specifically recommends resource-based constrained delegation (RBCD) because it is configured on the resource (Backup1) by setting msDS-AllowedToActOnBehalfOfOtherIdentity to allow the front-end computer account (Server2$) to delegate to the service SPN on the resource. This model follows least privilege, avoids broad domain-wide delegation, works with Protected Users, and requires minimal administration because you touch only the destination resource. In short: configure RBCD on Backup1 to permit Server2 to obtain S4U2Proxy tickets to the share service, enabling User1's PowerShell session on Server2 to access Backup1 while remaining compliant with Protected Users restrictions and Kerberos-only authentication.

NEW QUESTION: 104
You have a server that runs Windows Server 2022 and has the network adapters shown in the following table.

Name	Vendor	Interface	Link speed	Remote Direct Memory Access (RDMA) support
LAN1	Vendor1	Network Adapter #1	1 Gbps	No
LAN2	Vendor2	Network Adapter #2	25 Gbps	Yes
LAN3	Vendor3	Network Adapter #3	25 Gbps	Yes
LAN4	Vendor4	Network Adapter #4	1 Gbps	No

You need to configure NIC learning for LAN2 and LAN3. The solution must support Dynamic Virtual Machine Multi-Queue (d.VMMQ).
What should you use?
A. Static teaming mode
B. Switch Embedded Teaming (SET)
C. load balancing and failover (LBFO)
D. LACP teaming mode
Answer: B (LEAVE A REPLY)

The networking section of the AZ-800 material specifies that Switch Embedded Teaming (SET) is the teaming technology for Hyper-V hosts introduced in Windows Server 2016 and enhanced in later versions.
The guide highlights that SET "integrates teaming into the Hyper-V virtual switch," and supports advanced offloads including RDMA, VMMQ/d.VMMQ, SR-IOV, and RSS when the NICs support them. Conversely, the legacy LBFO (Load Balancing/Failover) teaming is described as "not supported with RDMA or VMMQ on Hyper-V hosts" and is deprecated for Hyper-V scenarios. Static or LACP teaming modes refer to LBFO team modes and thus inherit those limitations. Because LAN2 and LAN3 are 25-Gbps adapters that support RDMA and the requirement explicitly calls for NIC learning with Dynamic VMMQ, the prescribed and supported solution is to create a SET team using those two adapters and bind it to the Hyper-V switch. The documentation emphasizes that SET "enables RDMA with SMB Direct and Dynamic VMMQ on teamed NICs," meeting the performance and capability requirements.

NEW QUESTION: 105
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 Server 2022.
You need to ensure that Contained will run on Server1. The solution must minimize administrative effort.
What should you do?
A. Modify the Docker file.
B. Start Contained in Hyper-V isolation mode.
C. Start Contained in process isolation mode.
D. Rebuild the base image for Contained.
Answer: (SHOW ANSWER)

NEW QUESTION: 106
Task 5
You need to ensure that a DHCP scope named scope! on SRV1 can service client requests.
Answer:
See the solution of this Task below
Explanation:
One possible solution to ensure that a DHCP scope named scope1 on SRV1 can service client requests is to activate the scope on the DHCP server. A scope must be activated before it can assign IP addresses to DHCP clients. To activate a DHCP scope on SRV1, perform the following steps:
On SRV1, open DNS Manager from the Administrative Tools menu or by typing dnsmgmt.msc in the Run box.
In the left pane, expand your DHCP server and click on IPv4.

In the right pane, right-click on the scope that you want to activate, such as scope1, and select Activate.
Wait for the scope to be activated. You can verify the activation status by checking the icon next to the scope name. A green arrow indicates that the scope is active, while a red arrow indicates that the scope is inactive.
Now, the DHCP scope named scope1 on SRV1 can service client requests and lease IP addresses to DHCP clients. You can test the DHCP service by using the ipconfig /renew command on a DHCP client computer that is connected to the same subnet as the scope.

NEW QUESTION: 107
You have an Azure subscription that contains the storage accounts shown in the following table.

Name	Location	File share	Blob container
storage1	West US	share1	None
storage2	West US	share2	container2
storage3	Central US	share3	None
storage4	Central US	share4	container4

In the West US Azure region, you create a storage sync service named SyncA.
You plan to create a sync group named GroupA.
What is the maximum number of cloud endpoints you can use with GroupA?
A. 1
B. 2
C. 3
D. 4
Answer: A (LEAVE A REPLY)
The Azure File Sync section of the Administering Windows Server Hybrid Core Infrastructure materials states that a sync group defines the topology for synchronization and "contains one cloud endpoint and one or more server endpoints." A cloud endpoint is an Azure file share associated with a Storage Sync Service. The guidance also notes that "the Storage Sync Service and the storage account (file share) must reside in the same Azure region" and that "a single Azure file share can be the cloud endpoint for only one sync group." In the scenario, SyncA is in

West US, so only file shares in West US (for example, storage1\share1 or storage2\share2) are eligible. However, regardless of how many eligible shares exist, the exam guide is explicit: each sync group supports a maximum of one cloud endpoint. Additional endpoints in the group must be server endpoints on Windows Server volumes. Therefore, the maximum number of cloud endpoints you can use with GroupA is 1, which directly reflects the product's architecture and the documented exam objective requirements.

NEW QUESTION: 108
You have a Group Policy Object (GPO) named GPO1 that contains user settings only.
You plan to apply GPO1 to a global security group named Group1.
You link GPO1 to the domain, and you remove all the permissions granted to the Authenticated Users group.
You need to configure permissions for GPOI to meet the following requirements:
GPO1 must apply only to the users in Group1.
The solution must use the principle of least privilege.
Which permissions should you grant to Group1 and the Domain Computers group? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

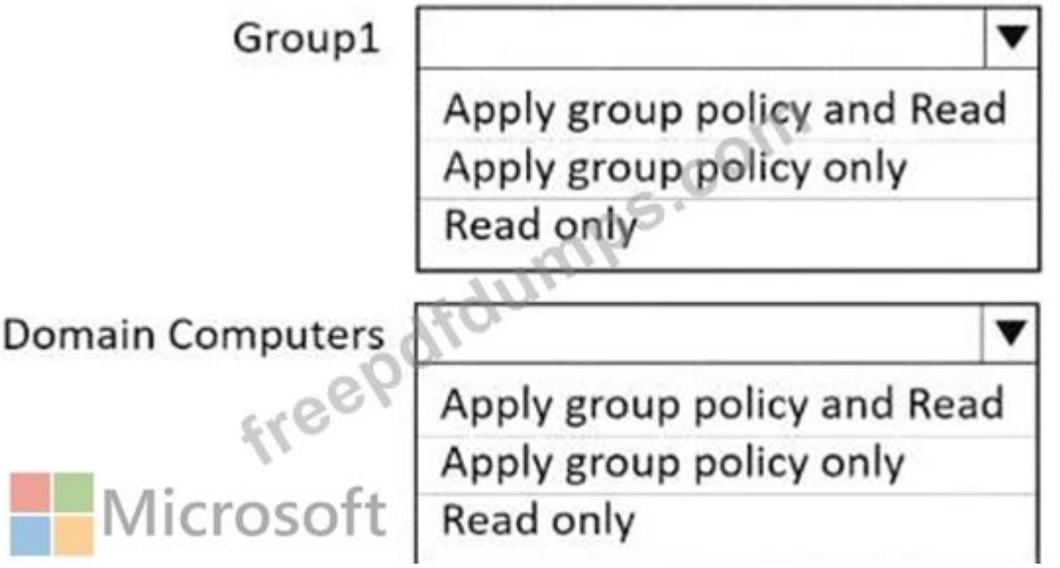

Answer:

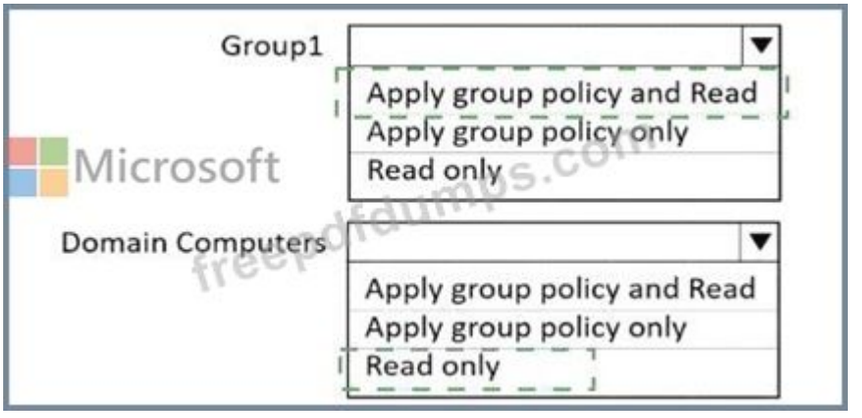

Explanation:

* Group1: Apply group policy and Read
* Domain Computers: Read only

In Windows Server Hybrid Core Infrastructure, the application of Group Policy Objects (GPOs) is controlled through security filtering. When you link a GPO to a domain or OU, it applies to all users and computers in that container by default because the Authenticated Users group is granted Read and Apply group policy permissions.

To meet the requirement that GPO1 applies only to users in Group1, you must first remove "Authenticated Users" from the security filter. To ensure the GPO actually applies to the target users, Group1 must be granted both Read and Apply group policy permissions. This allows the client-side extension (CSE) to read the policy settings and apply them to the user session.

Crucially, following security updates released by Microsoft (MS16-072), user-targeted GPOs are retrieved using the computer's security context rather than the user's. This change was implemented to prevent "man-in- the-middle" attacks. Therefore, for a user GPO to be processed, the computer from which the user is logging in must have Read access to the GPO. To adhere to the principle of least privilege, the Domain Computers group should be granted Read only access. Granting "Apply group policy" to Domain Computers is unnecessary and violates least privilege because the GPO only contains user settings and should not be processed in the computer's context beyond the initial retrieval.

NEW QUESTION: 109

Task2

You need to ensure that the Azure file share named share1 can sync to on-premises servers. The required source files are located in a folder named \\dc1.contoso.com\install.

You do NOT need to specify the on-premises servers at this time.

Answer:

See the solution of this Task below.

Explanation:

One possible solution to ensure that the Azure file share named share1 can sync to on-premises servers is to use Azure File Sync. Azure File Sync allows you to centralize your file shares in Azure Files without giving up the flexibility, performance, and compatibility of an on-premises file server. It does this by transforming your Windows Servers into a quick cache of your Azure file share. You can use any protocol available on Windows Server to access your data locally (including SMB, NFS, and FTPS) and you can have as many caches as you need across the world1.

Here are the steps to configure Azure File Sync for the Azure file share named share1 and the source files located in a folder named \dc1.contoso.com\install:

On the Azure portal, create a Storage Sync Service in the same region as your storage account that contains the Azure file share named share1. For more information on how to create a Storage Sync Service, see How to deploy Azure File Sync.

On the on-premises server that hosts the folder named \dc1.contoso.com\install, install the Azure File Sync agent. For more information on how to install the Azure File Sync agent, see Install the Azure File Sync agent.
On the on-premises server, register the server with the Storage Sync Service that you created in the first step.
For more information on how to register a server with a Storage Sync Service, see Register/unregister a server with Storage Sync Service.
On the Azure portal, create a sync group that defines the sync topology for a set of files. In the sync group, select the Azure file share named share1 as the cloud endpoint and the folder named \dc1.contoso.com\install as the server endpoint. For more information on how to create a sync group, see Create a sync group and a cloud endpoint and Create a server endpoint.
Wait for the initial sync to complete. You can monitor the sync progress on the Azure portal or on the on- premises server. For more information on how to monitor sync progress, see [Monitor sync progress].
Once the initial sync is complete, you can add more on-premises servers to the same sync group to sync and cache the content locally. You can also enable cloud tiering to optimize the storage space on the on-premises servers by tiering infrequently accessed or older files to Azure Files.

NEW QUESTION: 110
Your network contains an Active Directory domain named contoso.com. The domain contains group managed service accounts (gMSAs). You have a server named Server1 that runs Windows Server and is in a workgroup. Server! hosts Windows containers.
You need to ensure that the Windows containers can authenticate to contoso.com.
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.

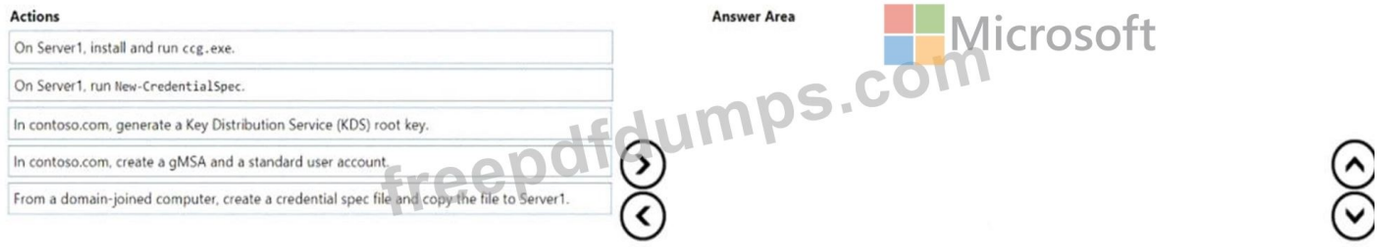

Answer:

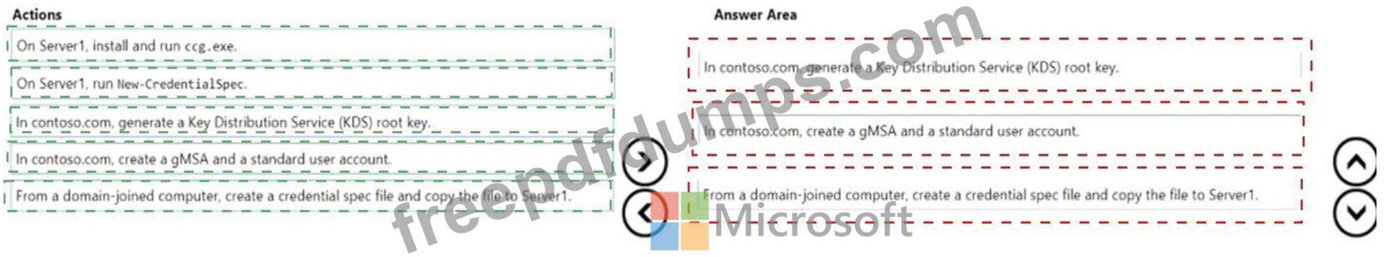

Explanation:

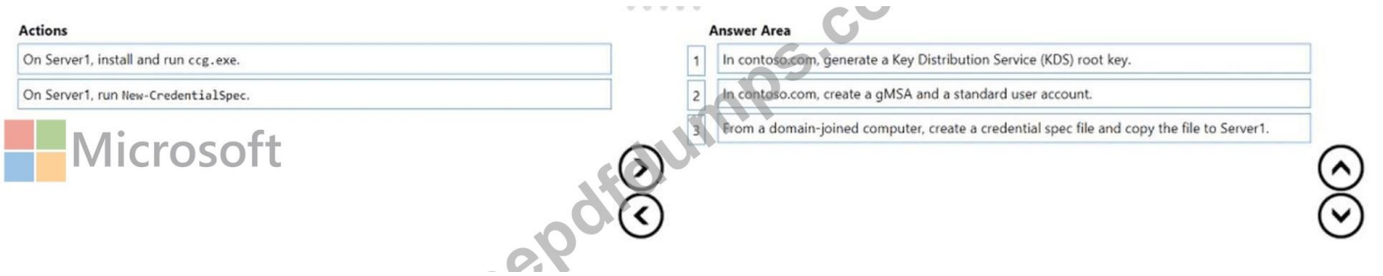

In Windows Server container scenarios, a containerized workload authenticates to Active Directory by using a group Managed Service Account (gMSA) referenced by a credential spec (CredSpec) JSON. The AZ-800 materials explain that gMSAs require the Key Distribution Service (KDS) root key once per forest before any gMSA passwords can be generated and rotated. Therefore, the first step is to create the KDS root key. Next, you create the gMSA and define who is allowed to retrieve its managed password. For workgroup (non- domain-joined) container hosts, you cannot authorize the host's computer account (because none exists), so you typically create a standard domain user (or group) and assign it in PrincipalsAllowedToRetrieveManagedPassword when creating the gMSA-this is called out in the hybrid core infrastructure guidance for container hosts. Finally, because Server1 is in a workgroup, you cannot run New-CredentialSpec on that host. The guidance states that the CredSpec file must be created on a domain- joined management machine (with AD tools), then copied to the workgroup host where containers are launched with --security-opt "credentialspec=...". Steps such as running New-CredentialSpec or other tooling directly on Server1-or granting broader rights than necessary-violate least-privilege and do not work in a workgroup configuration.

NEW QUESTION: 111

Task 8

You need to deploy a new primary DNS zone named fabrikam.com to DC1. The zone must be signed.

Answer:

See the solution of this Task below.

Explanation:

To deploy a new primary DNS zone named fabrikam.com to DC1 and sign the zone, you can follow these steps:

Step 1: Create the Primary DNS Zone Use the Add-DnsServerPrimaryZone PowerShell command to create the primary zone:

Add-DnsServerPrimaryZone -Name "fabrikam.com" -ZoneFile "fabrikam.com.dns" -DynamicUpdate Secure This command creates a primary zone for fabrikam.com with a DNS file named fabrikam.com.dns and allows secure dynamic updates.

Step 2: Sign the Zone To sign the zone, you can use the DNS Manager or Windows PowerShell. Here's how to sign the zone using PowerShell:

Add-DnsServerSigningKey -ZoneName "fabrikam.com" -Type KeySigningKey -CryptoAlgorithm RsaSha256 Set-DnsServerDnsSecZoneSetting -ZoneName "fabrikam.com" -DenialOfExistence

NSEC3 - NSEC3Parameters 1,0,10,"" These commands add a signing key to the zone and set DNSSEC settings with NSEC3 parameters.
Step 3: Publish the Signed Zone After signing the zone, ensure that it is published and available for DNS queries. You can verify the zone signing status using the following command:
Get-DnsServerZone -Name "fabrikam.com"
Note: Ensure that you have the appropriate permissions to perform these actions on DC1 and that the DNS Server role is installed and properly configured. Also, replace "fabrikam.com.dns" with the actual path to your DNS file if it's different12.
By following these steps, you should be able to deploy and sign the new primary DNS zone fabrikam.com on DC1.

NEW QUESTION: 112
You have an Azure virtual machine named VM1 that has a private IP address only.
You configure the Windows Admin Center extension on VM1.
You have an on-premises computer that runs Windows 11. You use the computer for server management.
You need to ensure that you can use Windows Admin Center from the Azure portal to manage VM1.
What should you configure?
A. an Azure Bastion host on the virtual network that contains VM1.
B. a VPN connection to the virtual network that contains VM1.
C. a network security group 1NSG) rule that allows inbound traffic on port 443.
D. a private endpoint on the virtual network that contains VM1.
Answer: B (LEAVE A REPLY)
According to the official documentation for Administering Windows Server Hybrid Core Infrastructure, using Windows Admin Center (WAC) within the Azure portal to manage virtual machines requires specific network pathing depending on the IP configuration of the target VM. For a virtual machine like VM1 that is configured with only a private IP address, the management traffic originates from the browser on the administrator's local computer and must be able to reach that private IP address directly.
The primary requirement for this scenario is that the management PC must have access to the virtual network (VNet) that is connected to the VM. As specified in the study guides, when managing a VM via its private IP, you must establish a hybrid connectivity solution such as a VPN connection (either Point-to-Site or Site-to-Site) or Azure ExpressRoute. This bridge allows the on-premises Windows 11 computer to route traffic into the Azure Virtual Network's private address space. Without this underlying network connectivity, the browser-based WAC gateway-which is hosted as an extension on the VM-cannot communicate with the portal interface.
While Azure Bastion (Option A) provides secure RDP and SSH access, the Windows Admin Center extension specifically requires network-level routing from the client to the private IP for its portal-integrated management features. Network Security Group (NSG) rules (Option C) are necessary to allow traffic on specific ports (like 443 for the WAC service), but they do not provide

the routing necessary to cross the boundary between an on-premises network and an Azure VNet. Similarly, Private Endpoints (Option D) are intended for accessing Azure PaaS services rather than providing a management path for IaaS VMs.
Therefore, configuring a VPN connection is the verified mandatory step to satisfy the connectivity requirements for WAC in this hybrid context.

NEW QUESTION: 113
Your network contains two Active Directory forests and a domain trust as shown in the following exhibit.

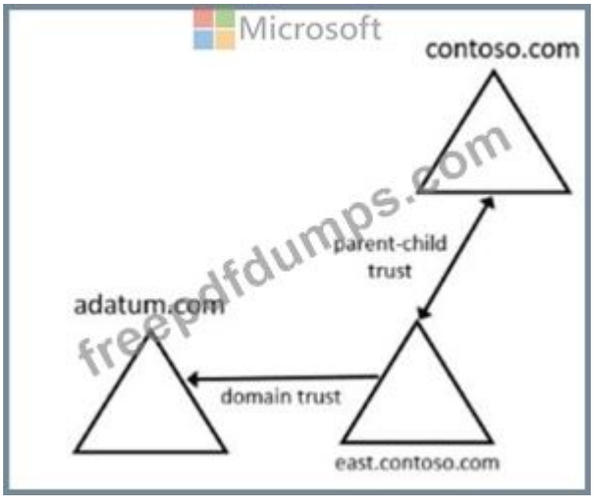

The domain trust has the following configurations:
* Name: adatum.com
* Type: External
* Direction: One-way. outgoing
* Outgoing trust authentication level: Domain-wide authentication

Name	Domain
User1	adatum.com
User2	contoso.com
User3	east.contoso.com

The forests contain the network shares shown in the following table.

Name	In domain
Share1	adatum.com
Share2	contoso.com
Share3	east.contoso.com

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

Answer:

Explanation:

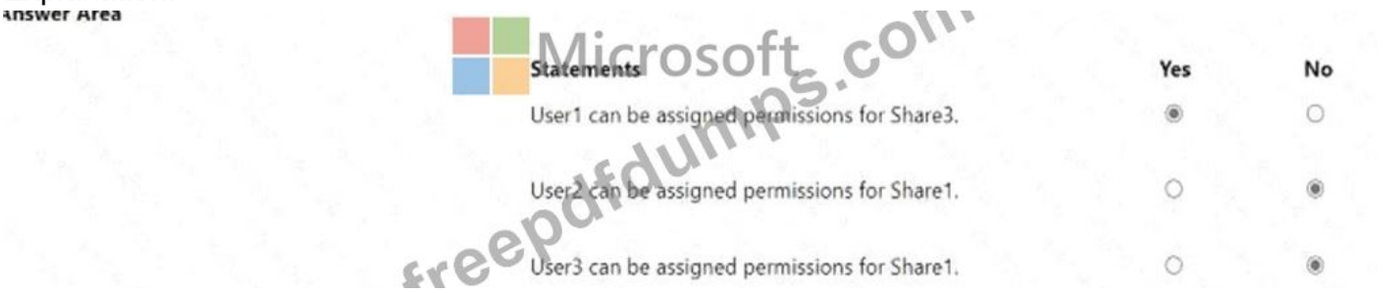

In the Administering Windows Server Hybrid Core Infrastructure guidance on trusts, Microsoft clarifies that a one-way outgoing trust from the local domain means the local domain trusts the specified external domain; therefore, "security principals from the trusted (specified) domain can be authenticated to access resources in the trusting (local) domain." Furthermore, external trusts are non-transitive and scoped only to the two domains that participate in the trust; they do not flow through to other domains or forests. The materials also state that parent-child trusts inside a forest are two-way and transitive, allowing authentication to cross between parent and child domains, but this transitivity does not extend across an external trust to another forest.

Applying these rules:

* The external one-way outgoing trust is configured on east.contoso.com # adatum.com. Thus, east.

contoso.com trusts adatum.com, allowing users from adatum.com to be granted access to resources in east.contoso.com. Therefore, User1 (adatum.com) can be assigned permissions on Share3 (east.contoso.

com) # Yes.

* User2 (contoso.com) attempting to access Share1 (adatum.com) relies on a path that would require trust from adatum.com to contoso.com or a transitive hop through east.contoso.com. Because the trust is external and non-transitive, and there's no trust from adatum.com to contoso.com, permissions cannot be assigned # No.

* User3 (east.contoso.com) to Share1 (adatum.com) would require adatum.com to trust east.contoso.com.

The configured direction is the opposite (east # adatum, outgoing), so east users are not trusted by adatum # No.

NEW QUESTION: 114

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.

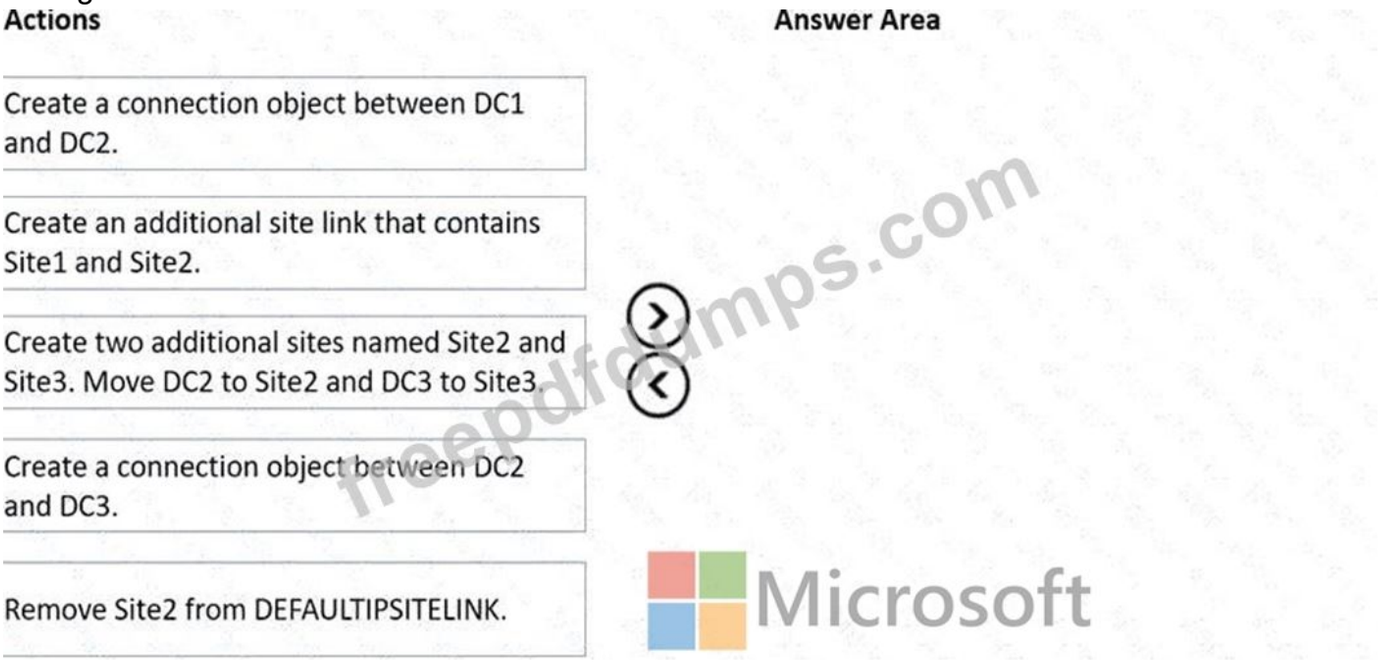

Answer:

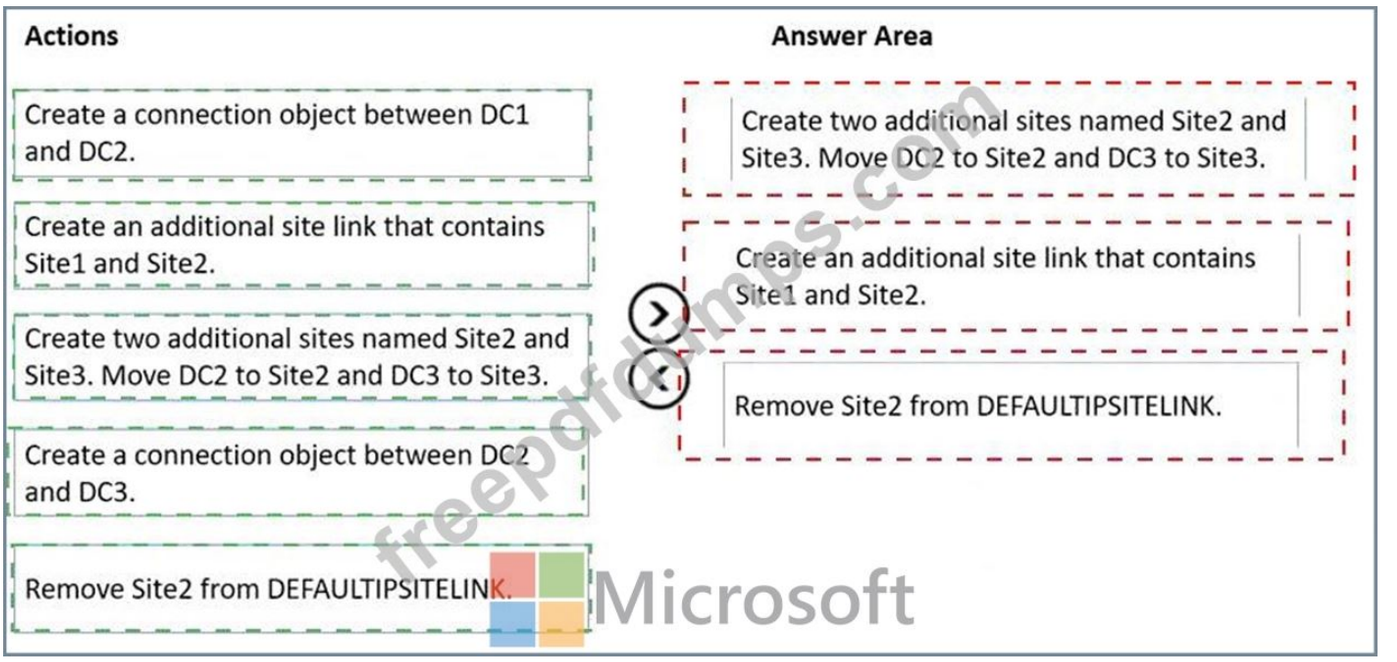

Explanation:

Create two additional sites named Site2 and Site3. Move DC2 to Site2 and DC3 to Site3.

Create an additional site link that contains Site1 and Site2.

Remove Site2 from DEFAULTIPSITELINK.

The Administering Windows Server Hybrid Core Infrastructure guidance for AD DS site design states that each physical location should be mapped to its own Active Directory site so that replication and logon traffic can be controlled and optimized per location. The documentation explains that when you deploy new sites, domain controllers must be moved into the corresponding site containers to ensure intrasite replication uses the local site topology and intersite replication follows site links. It also emphasizes that site links define replication cost, schedule, and frequency between sites; by creating separate site links for the pairs you want to manage, you can independently control replication schedules for those paths. The guides further note that new sites are placed in DEFAULTIPSITELINK by default, and that administrators can create additional site links and then remove a site from DEFAULTIPSITELINK to prevent it from inheriting the default schedule and cost-thereby isolating schedules per site pair. Finally, the KCC/ISTG automatically generate the required connection objects, so you generally should not create manual connection objects unless you have a specific exception. Following these principles: create Site2 and Site3 and move DCs, add a dedicated link (Site1-Site2) to set its own schedule, and remove Site2 from DEFAULTIPSITELINK so Site1-Site2 and Site1-Site3 replication can be controlled independently while minimizing replication interruptions.

NEW QUESTION: 115

You have an Azure subscription. The subscription contains a virtual machine named VM1 that runs Windows Server. You plan to manage VM1 by using a PowerShell runbook.

You need to create the runbook. What should you create first?

A. an Azure workbook

B. a Microsoft Power Automate flow

C. a Log Analytics workspace

D. an Azure Automation account

Answer: D (LEAVE A REPLY)

In the Administering Windows Server Hybrid Core Infrastructure objectives for automating management of Windows Server IaaS VMs, Microsoft specifies that runbooks reside in and are executed from an Azure Automation account. The Automation account is the management boundary that contains your runbooks, modules (Az PowerShell), credentials/variables, schedules, and identities. The guide explains that before you can create or publish a PowerShell or PowerShell 7 runbook, an Automation account must exist in the target subscription and region. After the account is created, you author the runbook, import any required Az modules, and grant permissions (commonly through the Automation account's managed identity) so the runbook can manage resources such as an Azure VM (VM1). While a Log Analytics workspace can be linked for job logs and update management, it is not required to create the runbook itself. Likewise, Power Automate is a separate service for workflow orchestration and Azure Workbooks are for monitoring/visualization; neither is the container for runbooks. Therefore, the first prerequisite to manage VM1 with a PowerShell runbook is to create an Azure Automation account, and then create the runbook within that account, assign permissions, and schedule or start it as needed.

NEW QUESTION: 116

Your network contains an Active Directory domain named contoso.com. The domain contains the computers shown in the following table.

Name	Operating system
Computer1	Windows 11
Server1	Windows Server 2016
Server2	Windows Server 2019
Server3	Windows Server 2022

On Server3, you create a Group Policy Object (GPO) named GP01 and link GPOI to contoso.com. GP01 includes a shortcut preference named Shortcut1 that has item-level targeting configured as shown in the following exhibit.

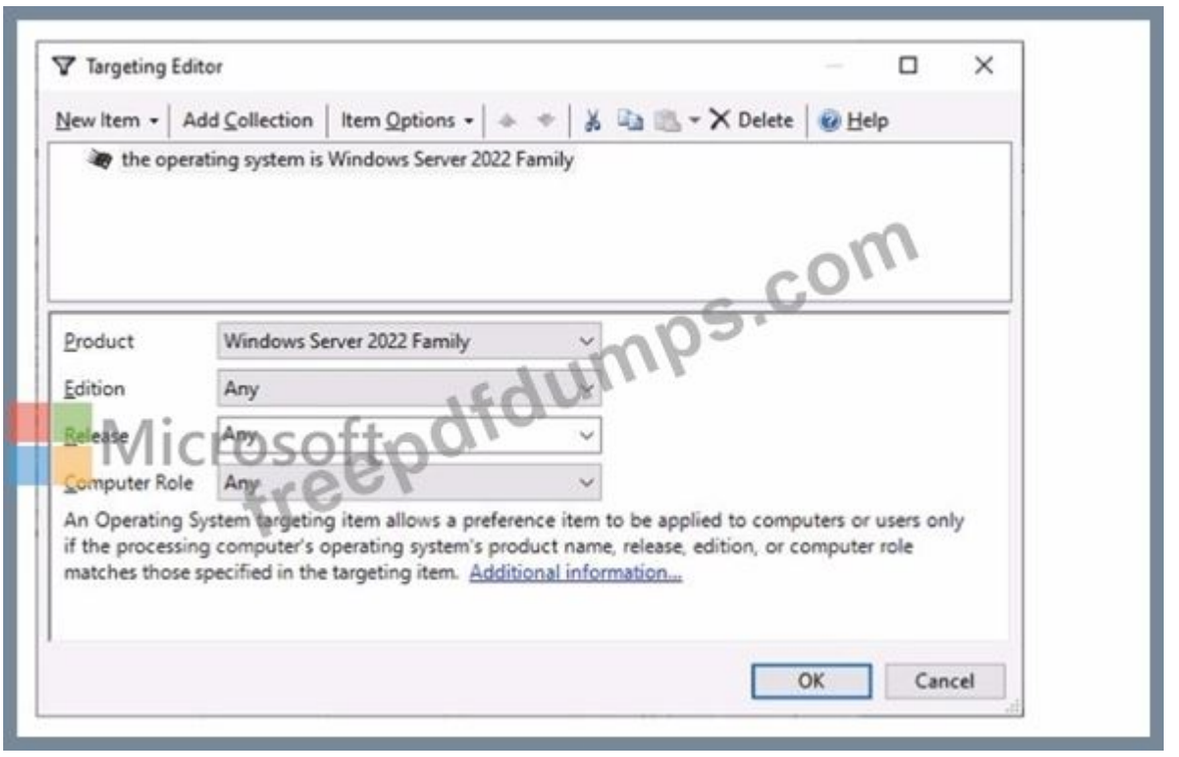

To which computer will Shortcut1 be applied?

A. Server3 only

B. Computer1 and Server3 only

C. Server2 and Server3 only

D. Server1, Server2, and Server3 only

Answer: A (LEAVE A REPLY)

Group Policy Preferences support Item-Level Targeting (ILT), allowing a preference item to apply only when the target computer meets specified criteria, such as operating system family and version. The AZ-800 study content notes that when a GPO is linked at the domain level, scope is all domain computers, but ILT on a specific preference item restricts that item to clients that match the ILT filter; non-matching clients still process the GPO but skip the filtered item. In the Targeting Editor shown, the condition is "the operating system is Windows Server 2022 Family." Among the listed machines: Computer1 (Windows 11), Server1 (Windows Server 2016), Server2 (Windows Server 2019), and Server3 (Windows Server 2022). Only Server3 satisfies the ILT. Therefore, the shortcut preference Shortcut1 is applied only to Server3.

NEW QUESTION: 117

You have an Azure virtual machine named Served that runs a network management application.

Server1 has the following network configurations:

* Network interface: Nic1
* IP address. 10.1.1.1/24
* Connected to: VnetVSubnet1

You need to connect Server1 to an additional subnet named Vnet1/Subnet2.

What should you do?

A. Modify the IP configurations of Nic1.

B. Add a network interface to Server1.
C. Add an IP configuration to Nic1.
D. Create a private endpoint on Subnet2
Answer: (SHOW ANSWER)
In Azure IaaS networking, a single network interface (NIC) can be attached to one subnet only. The AZ-
800 "Administering Windows Server Hybrid Core Infrastructure" study guidance for managing Windows Server IaaS VMs explains that NICs define the VM's Layer-3 placement and that "a NIC is associated with exactly one subnet; multiple IP configurations on a NIC must use addresses from the same subnet." Consequently, simply modifying the IP configuration of Nic1 (Option A) or adding a secondary IP configuration to Nic1 (Option C) will not move that interface into a second subnet; it only allows additional private IPs from the same subnet as the NIC's primary configuration.
To place the VM on Vnet1/Subnet2 in addition to Vnet1/Subnet1, you must multihome the VM by adding a second NIC that is attached to Vnet1/Subnet2 (Option B). The AZ-800 materials note that multinic support depends on the VM size, and that after adding the NIC you may need to adjust Windows networking (routing, binding order, firewall rules) to ensure the network management application listens on the correct interface(s). Private endpoints (Option D) are for securely exposing PaaS resources into a subnet and do not connect a VM to another subnet. Therefore, the only correct method to connect Server1 to an additional subnet is to add another network interface and attach it to Subnet2.

NEW QUESTION: 118
Your network contains an Active Directory Domain Services (AD DS) forest. The forest contains two domains named contoso.com and east.contoso.com. Contoso.com contains two users named CONTOSO
\User1 and EAST\User2.
You need to ensure that the users can perform the following tasks:
* User1 must deploy an additional domain controller to eastcontoso.com.
* User2 must deploy a new domain controller that will host a domain named west.contoso.com.
The solution must follow the principle of least privilege.
To which group should you add each user? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

The Administering Windows Server Hybrid Core Infrastructure guidance for AD DS promotion and forest operations states that adding a new domain controller to an existing domain requires credentials that "are a member of Domain Admins in the target domain (or equivalent delegated rights)." This is the minimum built-in role permitted to run AD DS installation and write to the domain's configuration containers for DC promotion. Therefore, to add an additional DC in east.contoso.com, the least-privilege group for User1 is EAST\Domain Admins.

For creating a new domain (child domain or new tree) in an existing forest, the exam materials specify that this is a forest-wide operation handled by the Domain Naming Master and requires enterprise-level permissions: "to create or remove domains in a forest, you must use an account that is a member of the Enterprise Admins group." Domain Admins in a single domain are insufficient because the task modifies forest-level naming contexts. Thus, to deploy the first DC for west.contoso.com, the least-privilege role that satisfies the requirement for User2 is CONTOSO\Enterprise Admins.

These selections follow the principle of least privilege: User1 is scoped to the child domain's administration only, while User2 receives the forest-level rights necessary to add a new domain.

NEW QUESTION: 119

You have a Windows Server container host named Server! that has a single disk. On Server1, you plan to start the containers shown in the following table.

Name	Description
Container1	Container1 is a Windows container that contains a web app in development. The container must NOT share a kernel with other containers.
Container2	Container2 is a Linux container that runs a web app. The container requires two static IP addresses.
Container3	Container3 is a Windows container that runs a database. The container requires a static IP address.

Which isolation mode can you use for each container? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Your network contains an Active Directory Domains Services (AD DS) domain named contoso.com. You implement a central store.

You create a new Group Policy Object (GPO) named GP01.

When you attempt to edit GP01, you see the settings shown in the exhibit. (Click the Exhibit tab.)

You need to ensure that all settings are available.

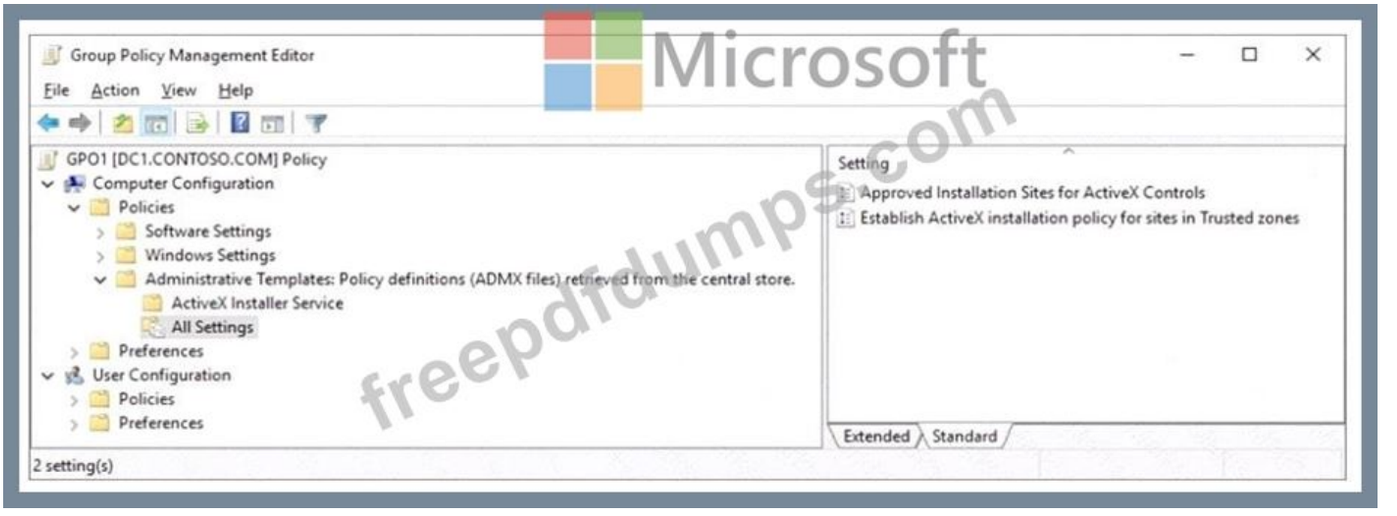

Solution: You delete the \\contoso.com\SYSVOL\contoso.com\Policies\PolicyDefinitions folder.

Does this meet the goal?

A. Yes

B. No

Answer: B (LEAVE A REPLY)

NEW QUESTION: 120

Task 4

You need to run a container that uses the mcrmicrosoft.com/windows/servercofe/iis image on SRV1. Port 80 on the container must be published to port 5001 on SRV1 and the container must run in the background See the solution of this Task below.

Answer:

To run a container on SRV1 using the mcrmicrosoft.com/windows/servercofe/iis image, publish port 80 on the container to port 5001 on SRV1, and ensure it runs in the background, you can follow these steps:
Step 1: Pull the IIS Image First, pull the correct IIS image from the Microsoft Container Registry:
docker pull mcr.microsoft.com/windows/servercore/iis
Step 2: Run the Container Next, run the container with the required port mapping and ensure it runs in the background using the -d flag:
docker run -d -p 5001:80 --name iis_container mcr.microsoft.com/windows/servercore/iis This command will start a container named iis_container using the IIS image, map port 80 inside the container to port 5001 on SRV1, and run the container in detached mode.
Step 3: Verify the Container is Running To verify that the container is running and the port is published, use the following command:
docker ps
This will list all running containers and show the port mappings.
Step
4: Access the IIS Server You can now access the IIS server running in the container by navigating to http://
<SRV1_IP>:5001 in a web browser, where <SRV1_IP> is the IP address of SRV1.
Note: Ensure that Docker is installed on SRV1 and that the port 5001 is open on the firewall to allow incoming connections1.
By following these steps, you should be able to run the IIS container on SRV1 with the specified port mapping and have it running in the background. Please replace mcrmicrosoft.com/windows/servercofe/iis with the correct image name mcr.microsoft.com/windows/servercore/iis as shown in the commands above.

NEW QUESTION: 121
You have a server named Server1 that runs Windows Server and contains a file share named Share1.
You need to prevent users from stoning MP4 files in Share1. The solution must ensure that the users can store other types of files in the share.
What should you configure on Server1?
A. File Management Tasks
B. NTFS Quotas
C. NTFS permissions
D. file screens
Answer: (SHOW ANSWER)
In the Administering Windows Server Hybrid Core Infrastructure content for File Services, Microsoft describes File Server Resource Manager (FSRM) as the feature used to control the types of files that users can store on file servers. The guide explains that a file screen "prevents users from saving unauthorized file types on a volume or in a folder by applying a file group (such as Audio and Video) or custom file name patterns (for example, *.mp4)." It further clarifies that file

screening "does not affect access permissions to other files and folders and allows permitted file types to be stored without restriction." By contrast, NTFS permissions govern access (read/write/modify) and "cannot filter by file extension," and NTFS quotas and File Management Tasks address capacity and automated tasks, not content blocking. Therefore, to stop users from storing .MP4 while allowing all other files, you create a File Screen on the path to Share1 and block either the built-in Audio and Video file group or a custom pattern *.mp4. This aligns precisely with the exam guide's directive that FSRM File Screening is the supported and intended method to restrict file types on Windows file shares while leaving other file operations unaffected.

NEW QUESTION: 122

You have a Windows Server 2022 container host named Host1 and a container registry that contains the container images shown in the following table.

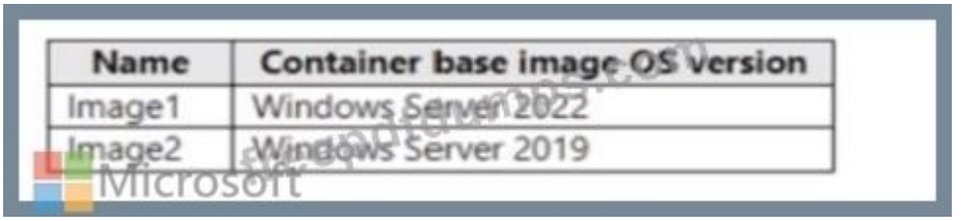

Name	Container base image OS version
Image1	Windows Server 2022
Image2	Windows Server 2019

You need to run the containers on Host1

Which isolation mode can you use for each image? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

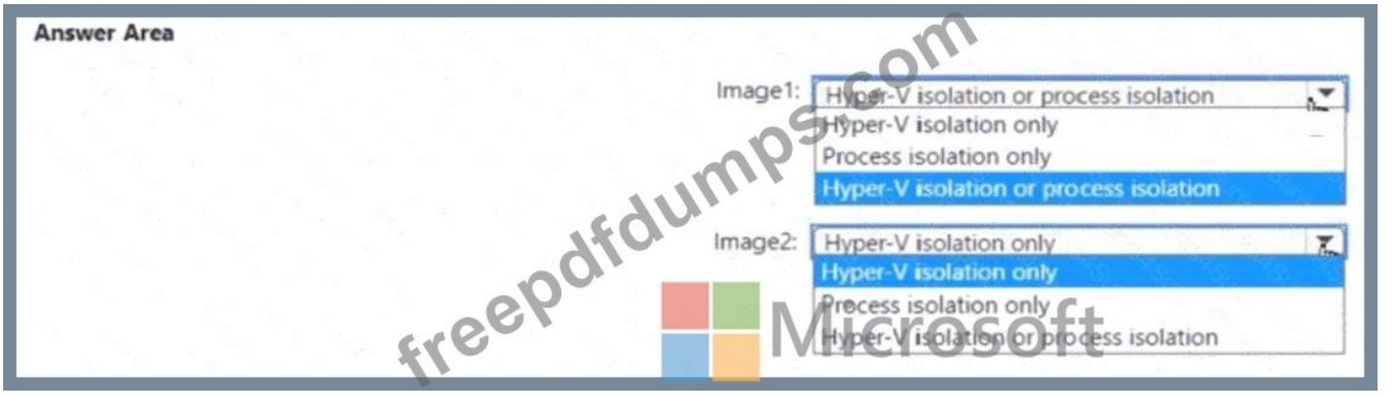

Answer:

Answer Area

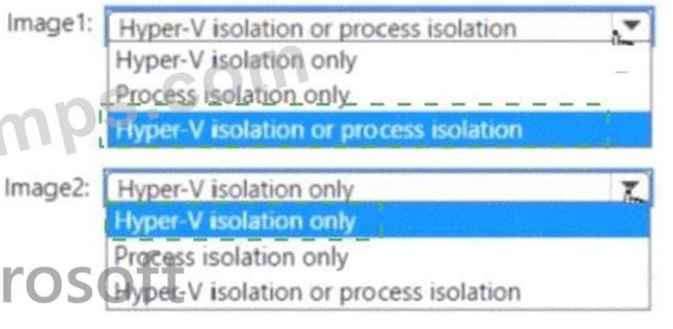

Explanation:

In the Windows Server container topic of Administering Windows Server Hybrid Core Infrastructure, Microsoft states the behavior of the two Windows container isolation modes very clearly:

* Process isolation: "Windows Server containers run with process isolation and share the kernel with the host. Because the container uses the host's kernel, the container base image must match the host OS version (including build) for the container to start."
* Hyper-V isolation: "With Hyper-V isolated containers, each container runs inside a lightweight virtual machine that has its own kernel, providing kernel-level isolation. Hyper-V isolation removes the requirement for the container image version to match the host, allowing older or different Windows versions to run on a newer host." Applying those rules to Host1 (Windows Server 2022):
* Image1 (Windows Server 2022) matches the host version. Therefore, it can run with process isolation (kernel shared and versions match). Hyper-V isolation is always available as well, so either process or Hyper-V isolation can be used.
* Image2 (Windows Server 2019) does not match the host's kernel version. With process isolation this mismatch prevents startup. However, Hyper-V isolation provides its own kernel boundary, so the container can run only with Hyper-V isolation.

Thus, the correct selections are: Image1 - Hyper-V isolation or process isolation; Image2 - Hyper-V isolation only.