DUMPSBOSS.

Upgrading Your Skills to MCSA: Windows Server 2016

Microsoft 70-743

Version Demo

Total Demo Questions: 25

Total Premium Questions: 258

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Install Windows Servers in host and compute environments	12
Topic 2, Implement storage solutions	26
Topic 3, Implement Hyper-V	68
Topic 4, Implement Windows containers	17
Topic 5, Implement high availability	22
Topic 6, Implement server environments	38
Topic 7, Implement a networking infrastructure	75
Total	258

QUESTION NO: 1

You have a server named Server1 that runs Windows Server 2016.

You plan to deploy Internet Information Services (IIS) in a Windows container.

You need to prepare Server1 for the planned deployment.

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

A. Install the Container feature.

B. Install Docker.

C. Install the Base Container Images.

D. Install the Web Server role.

E. Install the Hyper-V server role.

ANSWER: A B C

Explanation:

Installing the Container feature, installing Docker, and installing the Base Container Images are the required preparation steps for hosting IIS in a Windows Server 2016 container. The Containers feature provides the Windows Server container-host capability, including the OS components required to create and run Windows Server containers. Docker supplies the container engine and command-line tooling used to retrieve images, create containers, configure ports, and start the IIS workload. On Windows Server 2016, Docker can be installed by using the DockerMsftProvider PowerShell provider.

A container is instantiated from an image rather than from a traditional operating-system installation. Therefore, the host must obtain the relevant Windows container image before it can create the IIS container. In practice, this is commonly the supported IIS image, which includes IIS and is built on a compatible Windows base image; Docker pulls the image and any required base-image layers from the container registry. Image and host operating-system versions must be compatible, an especially important consideration for Windows Server 2016 container deployments. Microsoft's Windows container setup guidance is available at Set up your Windows container environment. For the IIS container image and usage examples, see Run your first Windows container.

QUESTION NO: 2

You are deploying a small network that has 30 client computers. The network uses the 192.168.1.0/24 address space. All computers obtain IP configurations from a DHCP server named Server1.

You install a server named Server2 that runs Windows Server 2016. Server2 has two network adapters named internal and Internet. Internet connects to an Internet service provider (ISP) and obtains the 131.107.0.10 IP address. Internal connects to the internal network and is configured to use the 192.168.1.250 IP address.

You need to provide Internet connectivity for the client computers.

What should you do?

A. On Server2, select the Internet and Internal network adapters and bridge the connections. From the DHCP console on Server1, authorize Server2.

B. On Server1, stop the DHCP server. On the Internal network adapter on Server 2, enable Internet Connection Sharing (ICS).

C. On Server2 run the New-NetNat -Name NAT1 -InternalIPInterfaceAddressPrefix 1 92.168.1.0/24 cmdlet. Configure Server1 to provide the 003 Router option of 1 31.107.0.10.

D. Install the Routing role service on Server2 and configure the NAT routing protocol. Configure Server1 to provide the 003 Router option of 192.168.1.250.

ANSWER: D

Explanation:

Install the Routing role service on Server2 and configure the NAT routing protocol. Configure Server1 to provide the 003 Router option of 192.168.1.250. This design makes Server2 the network address translation gateway between the private 192.168.1.0/24 LAN and the ISP-facing network. NAT translates traffic sourced from the private client addresses to Server2's public-facing address, 131.107.0.10, allowing client connections to reach the Internet while keeping the internal address space private.

Because Server2's Internal adapter is assigned 192.168.1.250, DHCP clients must receive that address as their default gateway. DHCP option 003 (Router) supplies this gateway setting to clients automatically. Clients then send nonlocal traffic to Server2, which applies NAT and routes it through the Internet adapter. The existing DHCP server can continue to assign addresses and related settings for the 30 clients, while Server2 provides the routing and translation function. Windows Server Routing and Remote Access supports NAT configuration for this multihomed-server scenario. See Microsoft Learn: NAT topology and Microsoft Learn: DHCP overview.

QUESTION NO: 3 - (HOTSPOT)

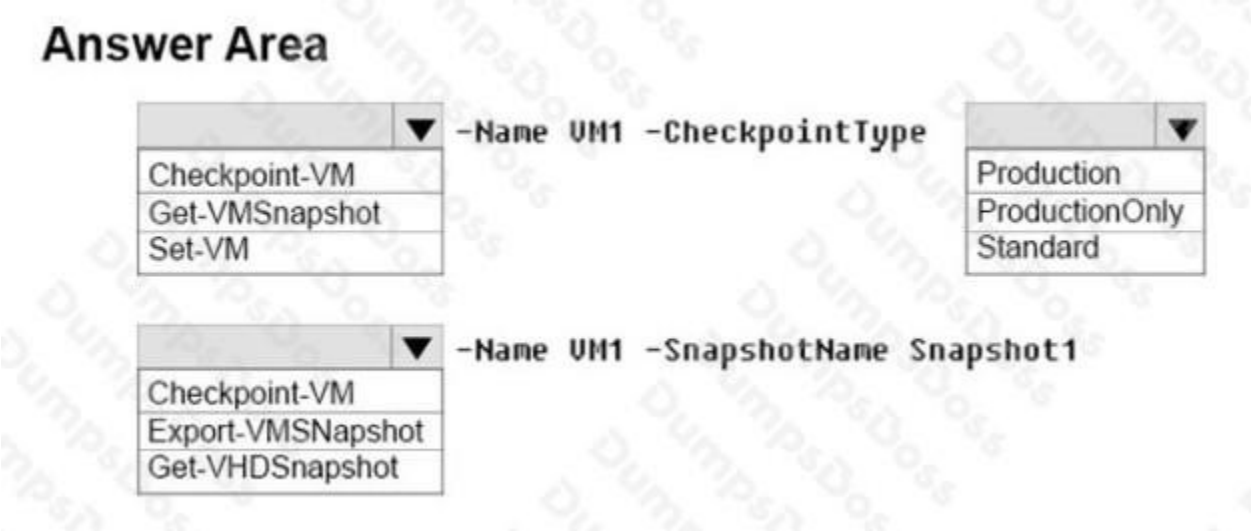
HOTSPOT

You have a server named VM1. VM1 is a virtual machine on a Hyper-V host that runs Windows Server 2016.

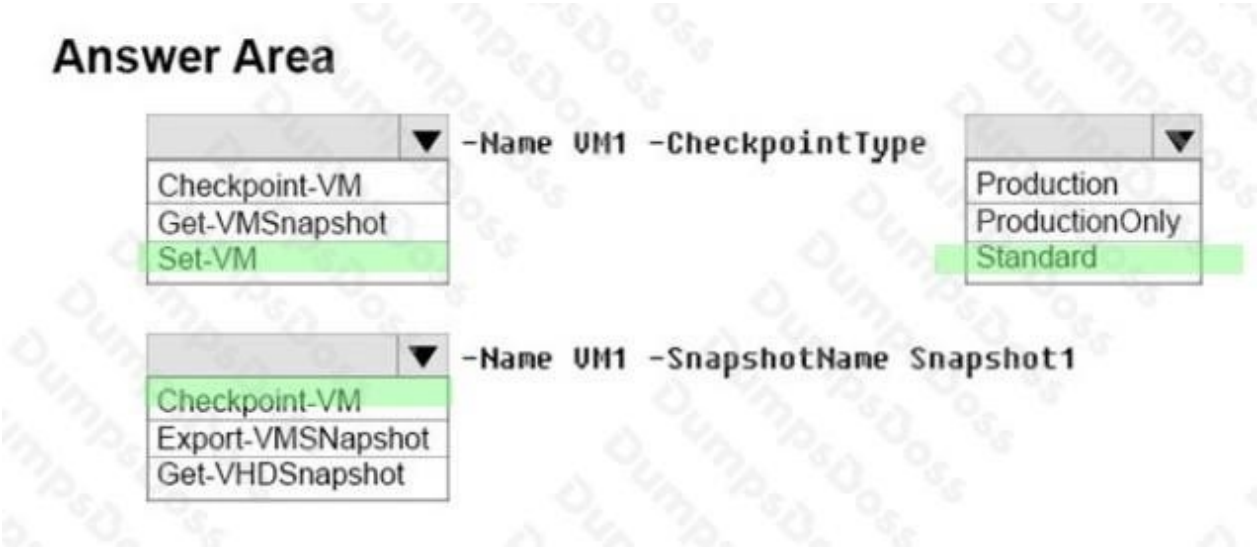
You need to create a checkpoint that includes the virtual machine memory state of VM1.

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

Hot Area:



ANSWER:



Explanation:

A standard checkpoint is the checkpoint type that includes the virtual machine's memory state. Before creating the checkpoint, the virtual machine must therefore be configured to use the Standard checkpoint type. The command `Set-VM -Name VM1 -CheckpointType Standard` applies that setting specifically to VM1.

After the checkpoint type is configured, `Checkpoint-VM -Name VM1 -SnapshotName Snapshot1` creates a checkpoint named Snapshot1 for that virtual machine. Because VM1 is configured for standard checkpoints, the created checkpoint captures the virtual machine's configuration, virtual disks, and the state of memory for a running VM. This allows the VM to be restored to the precise running condition present at the time the checkpoint was taken, rather than only restoring application-consistent disk data.

This behavior is particularly useful for development, testing, and troubleshooting situations where returning to an exact machine state matters. Hyper-V documents that standard checkpoints capture the state, data, and hardware configuration of a virtual machine, including saved memory state. See Microsoft's guidance on choosing between standard and production checkpoints and the PowerShell documentation for Set-VM and Checkpoint-VM.

QUESTION NO: 4

Your network contains an Active Directory domain. The domain contains two Hyper-V hosts.

You plan to perform live migrations between the hosts.

You need to ensure that the live migration traffic is authenticated by using Kerberos.

What should you do first?

A. From Server Manager, install the Host Guardian Service server role on a domain controller.

B. From Active Directory Users and Computers, add the computer accounts for both servers to the Cryptographic Operators group.

C. From Active Directory Users and Computers, modify the Delegation properties of the computer accounts for both servers.

D. From Server Manager, install the Host Guardian Service server role on both servers.

ANSWER: C

Explanation:

To authenticate Hyper-V live migration traffic with Kerberos, first configure constrained delegation on the Active Directory computer accounts for both Hyper-V hosts. Kerberos authentication requires delegation because a live migration initiated by an administrator must allow the source host to delegate that administrator's credentials to the destination host for the migration service. In Active Directory Users and Computers, open each host computer account's Delegation tab and

configure it to trust the computer for delegation to specified services only, using Kerberos only. Then add the required migration service entries for the other host, including the Microsoft Virtual System Migration Service. This configuration must be reciprocal because either server can be the source or destination of a live migration. After delegation is configured, set the live migration authentication protocol on the Hyper-V hosts to Kerberos and configure the permitted incoming live migration networks as appropriate. Microsoft documents constrained delegation as the prerequisite configuration for Kerberos-authenticated live migration between nonclustered hosts. See Set up hosts for live migration without Failover Clustering and Hyper-V live migration overview.

QUESTION NO: 5

You have a Remote Desktop Gateway (RD Gateway) server named Server1.

You need to ensure that members of a group named Contoso\Remote Users can connect through Server1 only to a server named AppServer1.

Which two policies should you configure? Each correct answer presents part of the solution.

A. A Remote Desktop Connection Authorization Policy (RD CAP)

B. A Remote Desktop Resource Authorization Policy (RD RAP)

C. A Network Policy Server connection request policy

D. An Active Directory Federation Services issuance authorization rule

E. A Windows Defender Firewall connection security rule

ANSWER: A B

Explanation:

A Remote Desktop Connection Authorization Policy (RD CAP) controls which users or groups are permitted to connect through an RD Gateway server. Configure an RD CAP that permits members of Contoso\Remote Users to connect through Server1.

A Remote Desktop Resource Authorization Policy (RD RAP) controls the internal network resources that authorized users can access through the RD Gateway server. Configure an RD RAP that permits access to AppServer1, either directly or through a computer group that contains AppServer1. Both policy types are required: an RD CAP authorizes the connection to the gateway, and an RD RAP authorizes the destination resource.

For more information, see Plan access from anywhere with Remote Desktop Services.

QUESTION NO: 6

You have a four-node failover cluster that runs Windows Server 2012 R2. The cluster hosts highly available virtual machines.

You plan to upgrade all the cluster nodes to Windows Server 2016 while maintaining workload availability.

You need to complete the cluster operating system rolling upgrade after the Windows Server 2016 nodes have been added to the cluster.

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

A. Add the Windows Server 2016 nodes to the cluster.

B. Drain the workloads from each Windows Server 2012 R2 node.

C. Remove each Windows Server 2012 R2 node from the cluster after it is drained.

D. Run Update-ClusterFunctionalLevel after all the nodes run Windows Server 2016.

E. Run Update-ClusterFunctionalLevel before adding the Windows Server 2016 nodes.

F. Enable Storage Spaces Direct before removing the Windows Server 2012 R2 nodes.

ANSWER: A B C D

Explanation:

A cluster operating system rolling upgrade allows Windows Server 2012 R2 and Windows Server 2016 nodes to coexist temporarily in the same cluster. Add the Windows Server 2016 nodes first. Then drain the workloads from each Windows Server 2012 R2 node and remove that node from the cluster. Repeat this process until all earlier-version nodes have been removed.

After every cluster node runs Windows Server 2016, run `Update-ClusterFunctionalLevel`. This raises the cluster functional level and enables Windows Server 2016 cluster features. Raising the functional level is irreversible, so it must be deferred until no earlier-version nodes remain in the cluster.

See Cluster operating system rolling upgrade and Update-ClusterFunctionalLevel.

QUESTION NO: 7 - (HOTSPOT)

HOTSPOT

You have two Hyper-V hosts named Server1 and Server2 that run Windows Server 2016. Server1 and Server2 connect to the same network.

Server1 and Server2 have virtual switches configured as shown in the following table.

Switch name	Host	Type	VLAN ID
Switch1	Server1	External	2
Switch2	Server2	External	4
Switch3	Server2	Internal	*Not applicable*

You have nine virtual machines configured as shown in the following table.

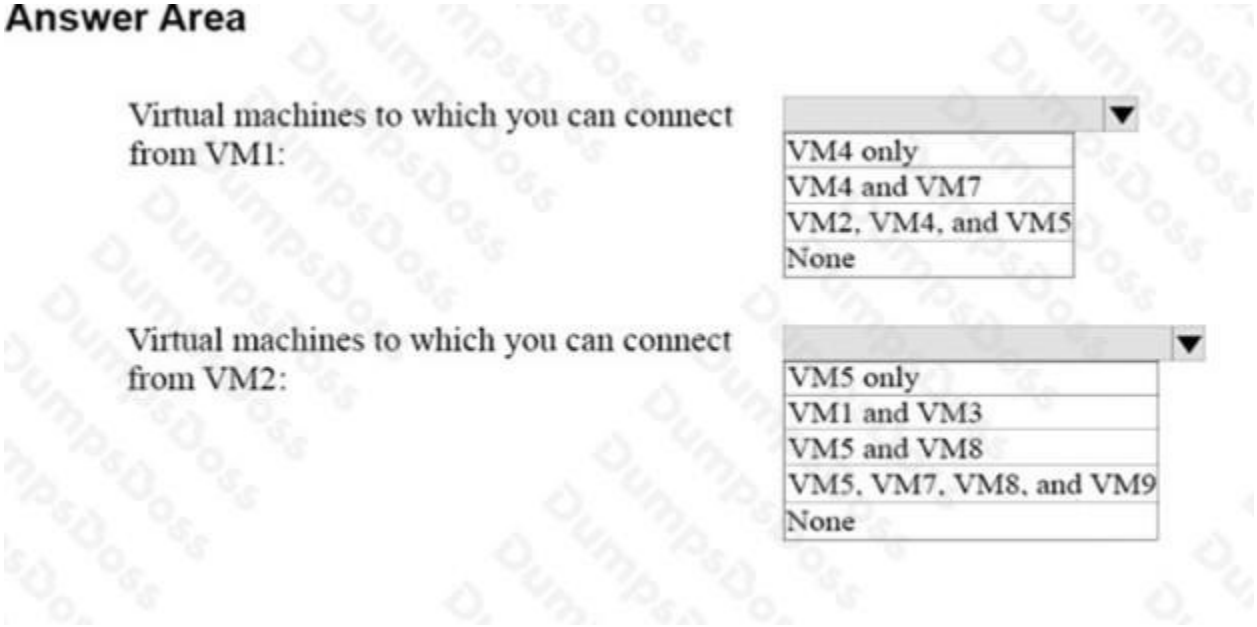
Virtual machine name	Connected to	VLAN ID
VM1	Switch1	2
VM2	Switch1	*Not applicable*
VM3	Switch1	4
VM4	Switch2	2
VM5	Switch2	*Not applicable*
VM6	Switch2	4
VM7	Switch3	2
VM8	Switch3	*Not applicable*
VM9	Switch3	4

All of the virtual machines are configured to have IP addresses from the same network segment. The firewall on each of the virtual machines is configured to allow network connectivity.

To which virtual machines can you connect from VM1 and VM2? To answer, select the appropriate options in the answer area.

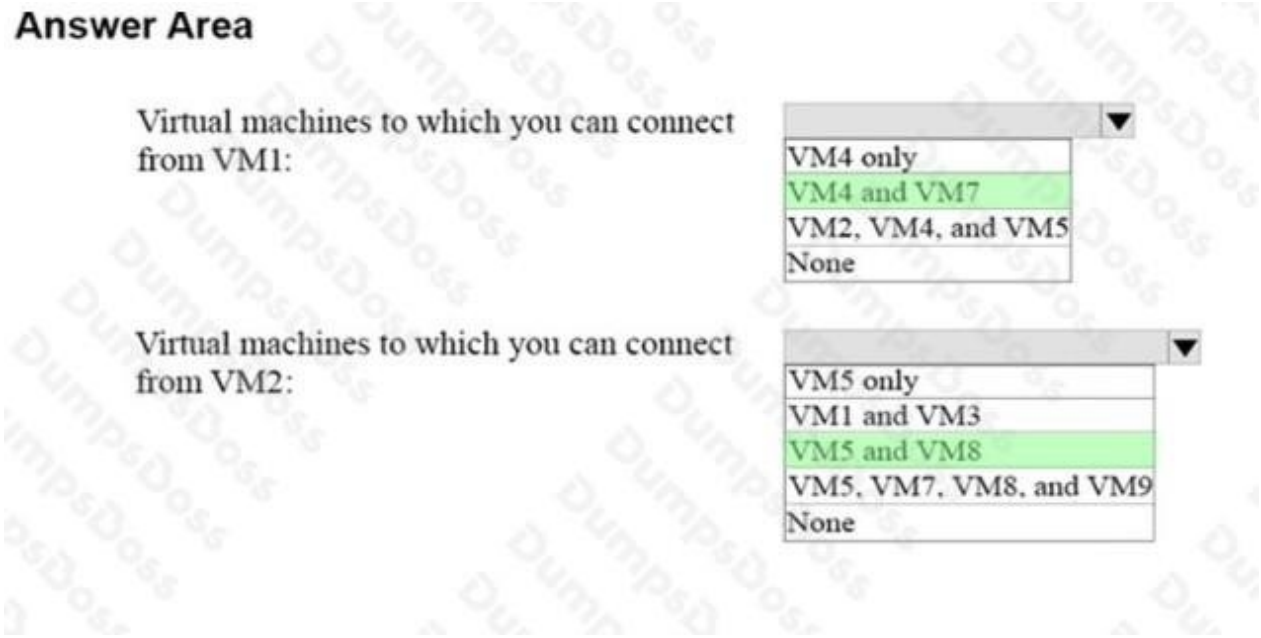
Hot Area:

Answer Area



ANSWER:

Answer Area



Explanation:

VM1 can connect to VM4 and VM7, and VM2 can connect to VM5 and VM8. Hyper-V virtual networking is determined by the virtual switch connection and the virtual LAN configuration associated with each virtual network adapter. An IP address in the same IP subnet does not, by itself, create connectivity: the virtual adapters must also be connected through the same Layer-2 broadcast domain. In this configuration, the virtual network used by VM1 includes VM4 and VM7, so those two virtual machines are reachable from VM1. The virtual network used by VM2 includes VM5 and VM8, so those two virtual machines are reachable from VM2.

Hyper-V switches provide the Layer-2 forwarding boundary for virtual machines. Where VLAN configuration is used, Hyper-V associates traffic with a VLAN and forwards frames only within the applicable VLAN path. This allows isolated virtual networks to use addressing from the same IP network segment without becoming mutually reachable. The relevant outcome here is that the switch and VLAN membership create one connectivity group containing VM1, VM4, and VM7, and another containing VM2, VM5, and VM8. Because the virtual-machine firewalls permit network traffic, no additional guest firewall restriction changes those switch-based reachability results.

For more detail, Microsoft explains Hyper-V virtual switch creation and its role in connecting virtual machines in Create a virtual switch for Hyper-V virtual machines. Microsoft also documents VLAN assignment for Hyper-V virtual network adapters

in Set-VMNetworkAdapterVlan. These mechanisms are why the correct results are based on the virtual switching topology rather than only on the IP subnet assigned inside each guest.

QUESTION NO: 8

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some questions sets might have more than one correct solutions, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a server named Server1 that runs Windows Server 2016. Server1 has the Hyper-V server role and Docker installed.

You pull the Microsoft/iis Docker image to Server1.

You need to view the available space in the microsoft/iis Docker image.

Solution: You run the command docker run –isolation hyperv -d microsoft/iis. You open Disk Management on Server1.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

The correct answer is **No**. Running `docker run --isolation=hyperv -d microsoft/iis` starts a container using Hyper-V isolation, but it does not expose the container image's filesystem capacity or free space through the host server's Disk Management console. Hyper-V isolation provides the container with a more strongly isolated execution environment by using a lightweight utility virtual machine; it does not make the image layers appear as a normally manageable disk in Windows Disk Management.

Docker images are composed of read-only layers managed by the Docker engine, while a running container receives an additional writable layer. These layers and their storage consumption are managed under Docker's data root rather than as a disk volume whose available space can be examined from Disk Management. To assess Docker disk usage, an administrator should use Docker tooling such as `docker system df`; to inspect free space from the container's perspective, the administrator can run an appropriate command inside the running container. Microsoft's Windows container isolation overview is available in Hyper-V isolation for Windows containers, and Docker documents storage reporting in docker system df.

QUESTION NO: 9 - (DRAG DROP)

DRAG DROP

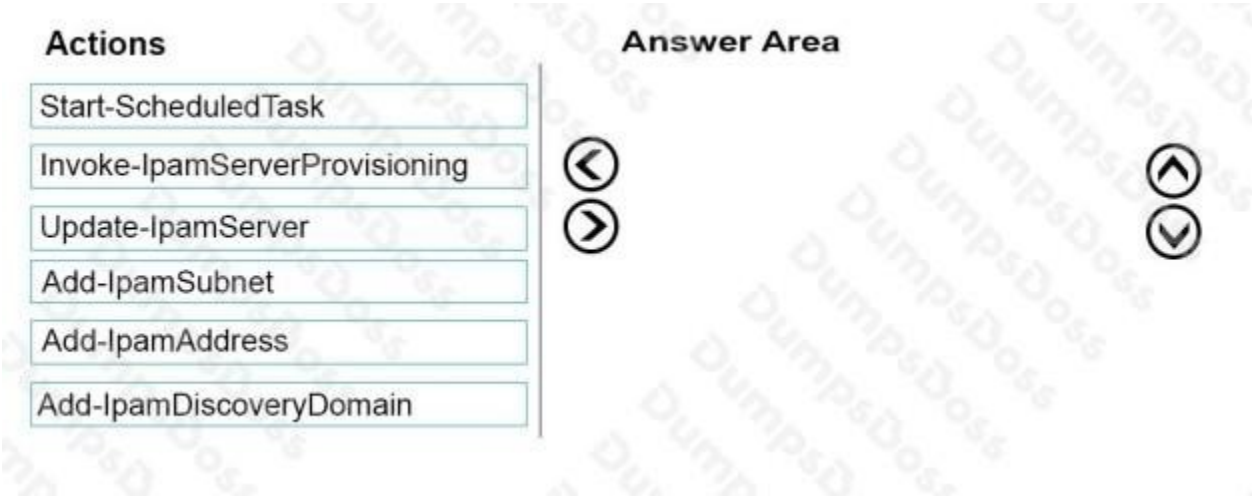
Your network contains an Active Directory domain named contoso.com. The domain contains a server named Server1 that runs Windows Server 2016.

You install IP Address Management (IPAM) on Server1.

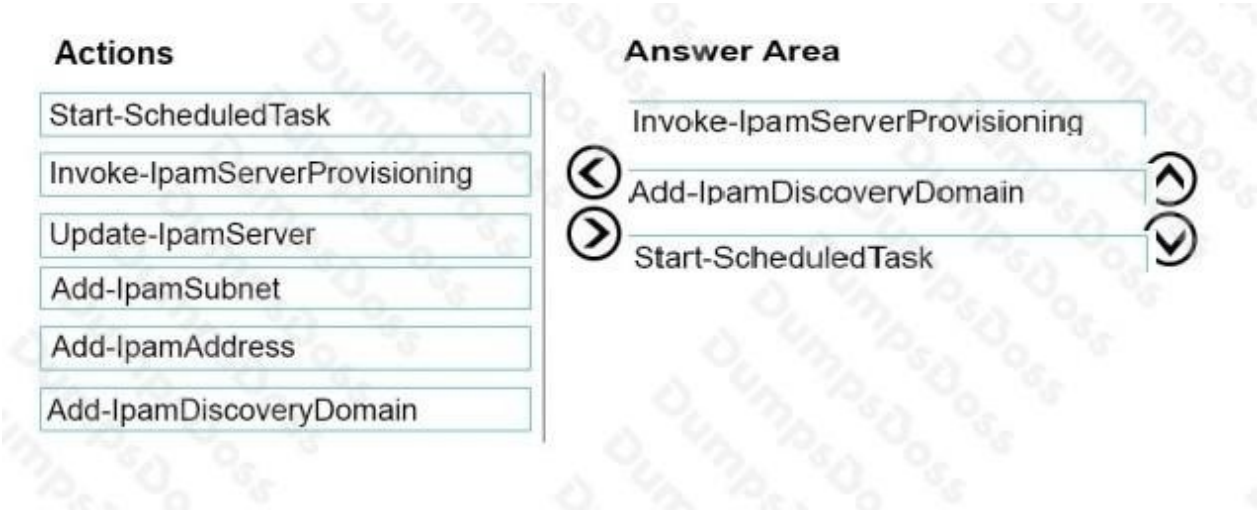
You need to manually start discovery of servers that IPAM can manage in contoso.com.

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

Select and Place:



ANSWER:



Explanation:

IPAM discovery must be configured before the server-discovery task is started. **Invoke-IpamServerProvisioning** is used first to establish the provisioning configuration for the IPAM server. Provisioning determines how IPAM obtains the permissions and access settings needed to manage the infrastructure servers it discovers. In a domain environment, this is the preparatory IPAM configuration step that enables the IPAM deployment to communicate with and manage the applicable server roles.

Next, **Add-IpamDiscoveryDomain** adds `contoso.com` to IPAM's discovery scope. A discovery domain tells IPAM where it should search Active Directory for manageable infrastructure servers. IPAM uses this configured scope to locate relevant domain controllers, DNS servers, and DHCP servers in the specified domain. This step is necessary before discovery begins because the IPAM server needs an explicit discovery domain configuration to know which domain to scan.

Finally, **Start-ScheduledTask** manually launches the IPAM ServerDiscovery scheduled task. This is the action that immediately initiates the discovery process rather than waiting for the task's normal schedule. Once the provisioning method and discovery-domain scope are in place, starting this task causes IPAM to query Active Directory and populate its inventory with discovered infrastructure servers. Microsoft documents the IPAM provisioning process in Invoke-IpamServerProvisioning and the discovery-domain configuration in Add-IpamDiscoveryDomain.

QUESTION NO: 10

Your network contains an Active Directory forest. The forest contains two domains named litwareinc.com and contoso.com. The contoso.com domain contains two domain controllers named LON-DC01 and LON-DC02. The domain controllers are located in a site named London that is associated to a subnet of 192.168.10.0/24.

You discover that LON-DC02 is not a global catalog server.

You need to configure LON-DC02 as a global catalog server.

What should you do?

A. From Active Directory Sites and Services, modify the NTDS Settings object of LON-DC02.

B. From Windows PowerShell, run the Set-NetNatGlobal cmdlet.

C. From Active Directory Sites and Services, modify the NTDS Settings object of the London site.

D. From the properties of the Domain Controllers organizational unit (OU) in Active Directory Users and Computers, modify the Security settings.

E. From the properties of the LON-DC02 computer account in Active Directory Users and Computers, modify the City attribute.

F. From Windows PowerShell, run the Enable-ADOptionalFeature cmdlet.

G. From the properties of the LON-DC02 computer account in Active Directory Users and Computers, modify the NTDS settings.

H. From Active Directory Sites and Services, modify the properties of the 192.168.10.0/24 IP subnet.

ANSWER: A

Explanation:

From Active Directory Sites and Services, modify the NTDS Settings object of LON-DC02. Global catalog status is configured on an individual domain controller's NTDS Settings object, not on the site, subnet, computer account, or Domain Controllers organizational unit. In Active Directory Sites and Services, expand the site containing the domain controller, expand the server object for LON-DC02, open the properties of its NTDS Settings object, and select the *Global Catalog* check box. After this setting is enabled, Active Directory replicates the required partial attribute set to that domain controller so it can provide forest-wide directory searches and participate in universal group membership processing where applicable.

A global catalog contains a searchable, read-only partial replica of objects from every domain in the forest, in addition to a full writable replica for the domain hosted by that domain controller. Microsoft documents both the role of the global catalog and its placement considerations at Global catalog server placement. The management procedure and the fact that the setting is exposed through the NTDS Settings properties are documented in The global catalog.

QUESTION NO: 11

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

Start of repeated scenario

Your network contains an Active Directory domain named contoso.com. The functional level of the forest and the domain is Windows Server 2008 R2. All servers in the domain run Windows Server 2016 Standard. The domain contains 300 client computers that run either Windows 8.1 or Windows 10.

The domain contains nine servers that are configured as shown in the following table.

Server name	Configuration	Planned changes
Server 1	Domain controller	None
Server 2	File server	Run Failover Clustering and Storage Spaces Direct
Server 3	File server	Run Failover Clustering and Storage Spaces Direct
Server 4	Hyper-V host	Run shielded virtual machines
Server 5	Hyper-V host	None
Server 6	Member server	Run Active Directory Federation Services (AD FS)
VM1	Virtual machine hosted on Server 5	None
VM2	Virtual machine hosted on Server 5	None
VM3	Virtual machine hosted on Server 5	None

The virtual machines are configured as follows:

- Each virtual machine has one virtual network adapter.
- VM1 and VM2 are part of a Network Load Balancing (NLB) cluster.
- All of the servers on the network can communicate with all of the virtual machines.

End of repeated scenario.

You plan to implement nested virtual machines on VM1.

Which two features will you be prevented from using for VM1?

A. NUMA spanning

B. Smart Paging

C. Dynamic Memory

D. live migration

ANSWER: C D

Explanation:

Implementing nested virtualization on VM1 requires Hyper-V to expose the physical processor virtualization extensions to the guest operating system. VM1 then acts as a Hyper-V host for the nested virtual machines. In the Windows Server 2016 implementation of nested virtualization, this configuration has explicit platform limitations for the outer virtual machine. Dynamic Memory cannot be used because nested virtualization requires a fixed amount of startup memory assigned to the virtualized Hyper-V host; Hyper-V cannot safely adjust that memory allocation while the VM is running. Live migration is also unavailable for a virtual machine configured for nested virtualization in this version. The processor virtualization state exposed to VM1 is tied to the physical host context, preventing VM1 from being moved live to another Hyper-V host. Therefore, VM1 must be configured with static memory, and operational planning must account for it remaining on its current host while nested virtualization is enabled. These limitations apply to the VM serving as the nested Hyper-V host, rather than

merely to the virtual machines created inside it. Microsoft documents nested-virtualization requirements and limitations at Nested virtualization and provides Windows Server 2016 Hyper-V feature guidance at Hyper-V overview.

QUESTION NO: 12

You have a file server named Server1 and a Hyper-V host named Server2 that run Windows Server 2016.

Server1 and Server2 each have two Remote Direct Memory Access (RDMA)-capable network adapters.

You need to maximize SMB file-transfer throughput and provide network-path resiliency for virtual machine files stored on Server1.

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

A. SMB Direct

B. SMB Multichannel

C. SMB Encryption

D. BranchCache

E. Data Deduplication

ANSWER: A B

Explanation:

SMB Direct uses RDMA-capable network adapters to provide low-latency, low-CPU-overhead SMB transfers. Because both servers have RDMA-capable adapters, SMB Direct can use those adapters for file access, including access to virtual machine files stored on an SMB share.

SMB Multichannel detects multiple available network paths between SMB clients and servers and can use them simultaneously. It increases throughput by using multiple connections and provides resiliency if one network path fails. SMB Multichannel and SMB Direct can operate together, allowing SMB to use multiple RDMA network connections.

SMB encryption protects data in transit but does not provide the required multiple-path throughput and resiliency. BranchCache is intended for caching content at branch offices. See SMB Direct and SMB Multichannel.

QUESTION NO: 13

You have an Active Directory domain that contains client computers that run Windows 10 Enterprise.

You plan to use BitLocker Network Unlock for operating system drives on the client computers.

You need to prepare the network infrastructure for Network Unlock.

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

A. Install the Windows Deployment Services server role.

B. Install the Network Unlock feature on the Windows Deployment Services server.

C. Enroll the Windows Deployment Services server for a Network Unlock certificate.

D. Install the Remote Access server role on the Windows Deployment Services server.

E. Configure a DHCP failover relationship for all DHCP scopes.

ANSWER: A B C

Explanation:

BitLocker Network Unlock requires a Windows Deployment Services (WDS) server with the Network Unlock feature installed. The WDS server provides the Network Unlock service that responds to supported clients during startup. The server must also have a Network Unlock certificate so that it can protect the network key material used in the startup process.

Clients require a TPM protector combined with a PIN protector and must be connected to a trusted wired network during startup. Network Unlock is not enabled merely by installing BitLocker on client computers or by configuring a normal DHCP scope. For more information, see BitLocker Network Unlock.

QUESTION NO: 14

Your network contains an Active Directory forest named contoso.com. The forest contains a member server named Server1 that runs Windows Server 2016. Server1 is located in the perimeter network.

You install the Active Directory Federation Services server role on Server1. You create an Active Directory Federation Services (AD FS) farm by using a certificate that has a subject name of sts.contoso.com.

You need to enable certificate authentication from the Internet on Server1.

Which two inbound TCP ports should you open on the firewall? Each correct answer presents part of the solution.

A. 389

B. 443

C. 3389

D. 8531

E. 49443

ANSWER: B E

Explanation:

Opening TCP 443 and TCP 49443 enables the required AD FS endpoints for Internet-facing certificate authentication. TCP 443 is the standard HTTPS port used by AD FS for federation service traffic, including access to the service name represented by the TLS certificate, sts.contoso.com. Internet clients use this protected HTTPS connection to reach AD FS.

TCP 49443 is the dedicated AD FS port for client TLS certificate authentication. AD FS uses this separate HTTPS listener for user certificate authentication because TLS client-certificate negotiation must occur during the TLS handshake, before the request can be processed by the normal AD FS application endpoints. Publishing and allowing this port from Internet clients allows AD FS to request and validate a client certificate as part of the authentication flow. Both ports must be reachable inbound through the perimeter firewall when AD FS itself is directly exposed in the perimeter network and certificate-based client authentication is required.

Microsoft documents that AD FS uses port 443 for HTTPS federation traffic and port 49443 for client certificate authentication. See AD FS requirements and Configure user certificate authentication in AD FS.

QUESTION NO: 15

You have a server named Server1 that runs Windows Server 2016 and has the Hyper-V server role installed.

On Server1, you plan to create a virtual machine named VM1.

You need to ensure that you can start VM1 from the network.

What are two possible ways to achieve the goal? Each correct answer presents a complete solution.

A. Create a generation 1 virtual machine that has a legacy network adapter.

B. Create a generation 1 virtual machine and run the Enable-NetAdapterPackageDirect cmdlet.

C. Create a generation 1 virtual machine and configure a single root I/O virtualization (SR-IOV) interface for the network adapter.

D. Create a generation 2 virtual machine.

ANSWER: A D

Explanation:

Creating a generation 1 virtual machine that has a legacy network adapter is a complete solution because the legacy network adapter emulates a traditional network interface and supports network boot. This adapter is specifically available for generation 1 virtual machines and can be used when booting from a network service is required. In the virtual machine firmware settings, the legacy network adapter can be placed in the appropriate startup order so the virtual machine attempts a network boot.

Creating a generation 2 virtual machine is also a complete solution. Generation 2 virtual machines use Unified Extensible Firmware Interface firmware and a synthetic network adapter rather than emulated legacy hardware. Hyper-V supports network boot for generation 2 virtual machines through that synthetic adapter, so no legacy adapter is required. The virtual machine's firmware boot order must be configured to prioritize the network adapter when network boot should occur before local boot devices. Microsoft documents the generation-specific virtual hardware and boot capabilities in its Hyper-V virtual machine generation planning guidance. For the network boot configuration process, see Hyper-V documentation.

QUESTION NO: 16

Your network contains an Active Directory domain named contoso.com. The domain contains a server named Server1 that runs Windows Server 2016.

Server1 has IP Address Management (IPAM) installed. IPAM is configured to use the Group Policy based provisioning method. The prefix for the IPAM Group Policy objects (GPOs) is IP.

From Group Policy Management, you manually rename the IPAM GPOs to have a prefix of IPAM.

You need to modify the GPO prefix used by IPAM.

What should you do?

A. Click Configure server discovery in Server Manager.

B. Run the Set-IpamConfiguration cmdlet.

C. Run the Invoke-IpamGpoProvisioning cmdlet.

D. Click Provision the IPAM server in Server Manager.

ANSWER: B

Explanation:

Run the Set-IpamConfiguration cmdlet. This cmdlet changes configuration settings on the IPAM server itself, including the GPO prefix that IPAM uses when it manages policy objects through Group Policy-based provisioning. After the policy objects were manually renamed from the IP prefix to the IPAM prefix, the IPAM configuration must be updated so that its stored prefix matches the renamed GPOs. Setting the GpoPrefix parameter to the new value aligns IPAM's configuration with the policy-object naming convention in the domain and enables IPAM to continue locating and using its provisioning policies correctly.

For Group Policy-based provisioning, IPAM uses GPOs to configure managed servers for tasks such as IP address tracking, DNS, and DHCP management. The GPO prefix is therefore a core IPAM provisioning setting, rather than merely a display label in Group Policy Management. Microsoft documents Set-IpamConfiguration as the cmdlet used to modify IPAM server configuration, including the GPO prefix. See Set-IpamConfiguration and IPAM deployment and provisioning.

QUESTION NO: 17

You have two servers named Server1 and Server2 that run Windows Server 2016.

You plan to implement Storage Replica to replicate the contents of volumes on Server1 to Server2.

You need to ensure that the replication traffic between the servers is limited to a maximum of 100 Mbps.

Which cmdlet should you run?

A. Set-NetUDPSetting

B. New-StorageQosPolicy

C. Set-SmbBandwidthLimit

D. Set-NetTCPSetting

E. Set-NetworkController

F. New-NetTransportFilter

G. Set-StorageQosPolicy

H. New-NetQosPolicy

ANSWER: C

Explanation:

`Set-SmbBandwidthLimit` is the correct cmdlet because Storage Replica transfers replication data by using SMB 3.x. Windows Server provides SMB bandwidth management categories so that administrators can apply throughput limits to specific types of SMB traffic, including the `StorageReplica` category. For this requirement, the cmdlet can be used with that category and an appropriate `-BytesPerSecond` value, such as `Set-SmbBandwidthLimit -Category StorageReplica -BytesPerSecond 12500000`. This value represents 100 megabits per second when using decimal network-rate units: 100,000,000 bits per second divided by eight equals 12,500,000 bytes per second.

Applying the limit at the SMB layer is appropriate because it directly governs the protocol carrying the replication stream, rather than attempting to control unrelated TCP, UDP, storage I/O, or general network settings. The configured bandwidth limit applies to SMB traffic matching the selected category and helps prevent Storage Replica replication from consuming more than the designated network capacity. Microsoft documents the available SMB bandwidth-limit categories and the `-BytesPerSecond` parameter in the Set-SmbBandwidthLimit cmdlet reference. Storage Replica's use of SMB is described in the Storage Replica overview.

QUESTION NO: 18

You have a container host named Server1 that runs Windows Server 2016.

You need to start a Hyper-V container on Server1.

Which parameter should you use with the docker run command?

A. –-runtime

B. --isolation

C. --entrypoint

D. --privileged

E. --expose

ANSWER: B

Explanation:

The correct parameter is `--isolation`. Windows containers support distinct isolation modes, and Docker selects the mode for an individual container by using the `--isolation` option with `docker run`. To explicitly start a Hyper-V isolated container, the command uses `--isolation=hyperv`, for example: `docker run --isolation=hyperv mcr.microsoft.com/windows/nanoserver`. Hyper-V isolation runs the container inside a lightweight utility virtual machine, rather than sharing the host kernel in the usual process-isolation model. This provides an additional isolation boundary and enables a container image whose Windows version is not directly compatible with the host to run in supported scenarios. On Windows Server 2016, specifying the isolation mode explicitly is the appropriate way to request a Hyper-V container at launch time. The setting is applied when the container is created by the `docker run` command; it is not an application startup, network publishing, or privilege setting. Microsoft documents Hyper-V isolation as one of the supported Windows container isolation modes and shows `--isolation=hyperv` as the Docker command-line syntax. See Hyper-V isolation for Windows containers and Windows container version compatibility.

QUESTION NO: 19

You have two servers named Server1 and Server2 that run Windows Server 2016. You plan to configure Storage Replica to replicate a data volume from Server1 to Server2.

You need to prepare the servers for the replication partnership.

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

A. Install the Storage Replica feature on Server1 and Server2.

B. Create a data volume on Server1 and Server2.

C. Create a dedicated log volume on Server1 and Server2.

D. Install the Failover Clustering feature on Server1 and Server2.

E. Configure both data volumes as Cluster Shared Volumes.

ANSWER: A B C

Explanation:

Storage Replica must be installed on both Server1 and Server2 because both servers participate in the block-level replication partnership. Each server also requires a dedicated data volume and a dedicated log volume. The data volumes contain the replicated application data, while the log volumes retain the write-order log used to maintain replication consistency.

The log volume must be separate from the data volume and should use fast storage because replication performance depends on log-write latency. After these requirements are met, a replication group and partnership can be created by using cmdlets such as `New-SRPartnership`. For more information, see Storage Replica overview and Server-to-server Storage Replica.

QUESTION NO: 20

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have a Hyper-V host named Server1 that runs Windows Server 2016. Server1 has a dynamically expanding virtual hard disk (VHD) file that is 900 GB. The VHD contains 400 GB of free space.

You need to reduce the amount of disk space used by the VHD.

What should you run?

A. the Mount-VHD cmdlet

B. the Diskpart command

C. the Set-VHD cmdlet

D. the Set-VM cmdlet

E. the Set-VMHost cmdlet

F. the Set-VMProcessor cmdlet

G. the Install-WindowsFeature cmdlet

H. the Optimize-VHD cmdlet

ANSWER: H

Explanation:

The **Optimize-VHD cmdlet** is the appropriate Hyper-V management command for reducing the physical storage consumed by a dynamically expanding VHD. A dynamically expanding virtual disk can grow as data is written to it, but deleting files inside the guest does not automatically return the corresponding blocks to the host file system. Optimizing the disk with the `-Mode Full` parameter compacts the VHD by reclaiming unused blocks and packing allocated blocks more efficiently, thereby reducing the VHD file's size on Server1.

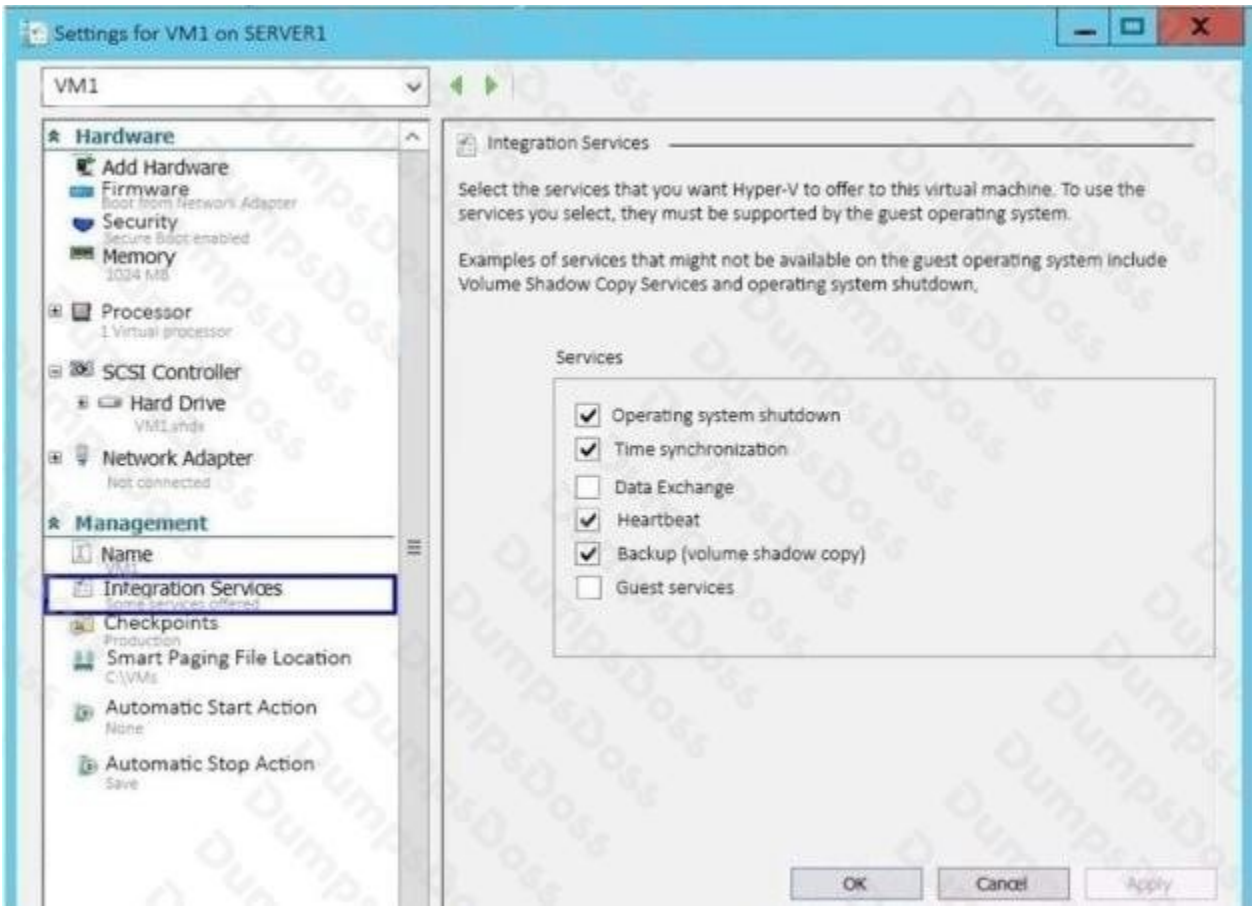
For a successful full compaction, the virtual disk must be detached or attached read-only. In practice, the guest operating system should first make its unused space reclaimable, such as by using an appropriate zeroing or trim/unmap process where supported; then the administrator can run a command such as `Optimize-VHD -Path "C:\VMs\Disk.vhd" -Mode Full`. This action is designed specifically for dynamic and differencing virtual hard disks and is not applicable to fixed-size disks. Microsoft documents that Optimize-VHD optimizes the allocation of space in VHD and VHDX files, including compacting dynamically expanding disks. See Optimize-VHD cmdlet documentation and Hyper-V virtual hard disk management.

QUESTION NO: 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 questions sets might have more than one correct solutions, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a Hyper-V host named Server1 that hosts a virtual machine named VM1. Server1 and VM1 run Windows Server 2016. The settings for VM1 are configured as shown in the exhibit below:



You need to ensure that you can use the Copy-VMFile cmdlet on Server1 to copy files from VM1. Solution: You start the Hyper-V Guest Service Interface service on VM1 Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct. The `Copy-VMFile` cmdlet is designed to copy files from the Hyper-V host into a guest virtual machine; it does not provide a guest-to-host file-copy mechanism. Therefore, starting the Hyper-V Guest Service Interface service in VM1 cannot enable copying files *from* VM1 to Server1 by using this cmdlet. The cmdlet's source path is a path accessible to the host, while its destination path is within the specified virtual machine.

For supported host-to-guest copying, Hyper-V Guest Services must also be enabled for the virtual machine in its Integration Services configuration. Guest Services supplies the host/guest integration channel used by `Copy-VMFile`, and the guest-side Hyper-V Guest Service Interface service participates in that channel. However, even when that service and integration component are available, the operation remains host-to-guest only. A different supported transfer method, such as an SMB share, PowerShell remoting, or an enhanced-session/remote-desktop transfer workflow, is needed when files must originate in VM1 and be copied to Server1.

Microsoft documents the cmdlet direction and requirements in Copy-VMFile and describes the Guest Services integration component in Manage Hyper-V Integration Services.

QUESTION NO: 22

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have a two-node Hyper-V cluster named Cluster1 at a primary location and a stand-alone Hyper-V host named Server1 at a secondary location.

A virtual machine named VM1 runs on Cluster1.

You configure a Hyper-V Replica of VM1 to Server1.

You need to perform a Test Failover of VM1.

Which tool should you use?

A. the clussvc.exe command

B. the cluster.exe command

C. the Computer Management snap-in

D. the configurehyperv.exe command

E. the Disk Management snap-in

F. the Failover Cluster Manager snap-in

G. the Hyper-V Manager snap-in

H. the Server Manager app

ANSWER: G

Explanation:

The **Hyper-V Manager snap-in** is the appropriate tool because the replica VM is hosted on the stand-alone secondary Hyper-V server, Server1. A Test Failover is initiated at the replica site and creates a separate temporary test virtual machine from a selected recovery point. This lets an administrator validate that the replicated VM can start and that its applications and configuration operate as expected, without interrupting ongoing replication or affecting the protected production VM on Cluster1.

In Hyper-V Manager, connect to Server1, select the replicated VM, and use the *Replication* menu to start *Test Failover*. Hyper-V prompts for the recovery point to use and starts the test VM with a distinct test configuration. After validation, use *Stop Test Failover* to remove the temporary test VM and return the replica to its normal state. Microsoft documents Test Failover as a replica-site operation and describes managing Hyper-V Replica through Hyper-V Manager. See Set up Hyper-V Replica and Manage Hyper-V.

QUESTION NO: 23

You have a server named Server1 that runs Windows Server 2016.

The disks on Server1 are configured as shown in the following table.

Volume	Type	File System	Capacity
C:	Attached locally	NTFS	150 GB
D:	Attached locally	exFAT	100 GB
E:	Attached locally	NTFS	20 GB
F:	Attached locally	ReFS	1 TB
G:	iSCSI LUN	NTFS	2 TB

Windows Server 2016 is installed in C:\Windows.

On which two volumes can you enable data deduplication? Each correct answer presents a complete solution.

A. C:

B. D:

C. E:

D. F:

E. G:

ANSWER: C E

Explanation:

E: and **G:** are the volumes that meet the Windows Server 2016 Data Deduplication prerequisites shown in the disk configuration. Data Deduplication is enabled per volume and requires a suitable data volume formatted with NTFS. Both E: and G: are eligible non-operating-system NTFS volumes in the configuration, so either can have the Data Deduplication role service enabled and then be configured with an appropriate usage type, such as General purpose file server or Virtual Desktop Infrastructure.

Data Deduplication works by identifying repeated chunks of data within a volume, retaining a single copy in the chunk store, and transparently replacing duplicate data with references to that shared content. Applications and users continue to access the files normally; the optimization occurs at the file-system level. In Windows Server 2016, this feature is intended for supported NTFS data volumes, including appropriately presented fixed storage volumes. Selecting E: and G: therefore provides two complete, independently valid locations for deduplication processing while leaving the Windows installation volume separate from the deduplicated data set.

See Data Deduplication overview and Install and enable Data Deduplication for the supported-volume requirements and configuration process.

QUESTION NO: 24

You have a Nano Server named Nano1.

You deploy several containers to Nano1 that use an image named Image1.

You need to deploy a new container to Nano1 that uses Image1.

What should you run?

A. the Install-WindowsFeature cmdlet

B. the docker run command

C. the docker load command

D. the Install-NanoServerPackage cmdlet

ANSWER: B

Explanation:

The **docker run command** is used to create and start a new container from a specified container image. In this case, Image1 is already present and has been used for several existing containers on Nano1. Running `docker run Image1`, optionally with a tag, command, port mappings, volumes, or other runtime parameters, instructs Docker to create a new container instance whose initial filesystem and configuration are based on that image. Each invocation creates a distinct container, allowing multiple independently running containers to share the same underlying image layers efficiently.

For example, `docker run Image1` starts a container using Image1's configured default command, while `docker run -d --name NewContainer Image1` creates a named container and runs it in the background. Docker uses a local copy of the image when it is available; otherwise, it attempts to retrieve the image from its configured registry. This is the standard Docker workflow for deploying an additional container from an existing image. See the Docker run command reference and Microsoft's Windows container quick-start guidance.

QUESTION NO: 25 - (DRAG DROP)

DRAG DROP

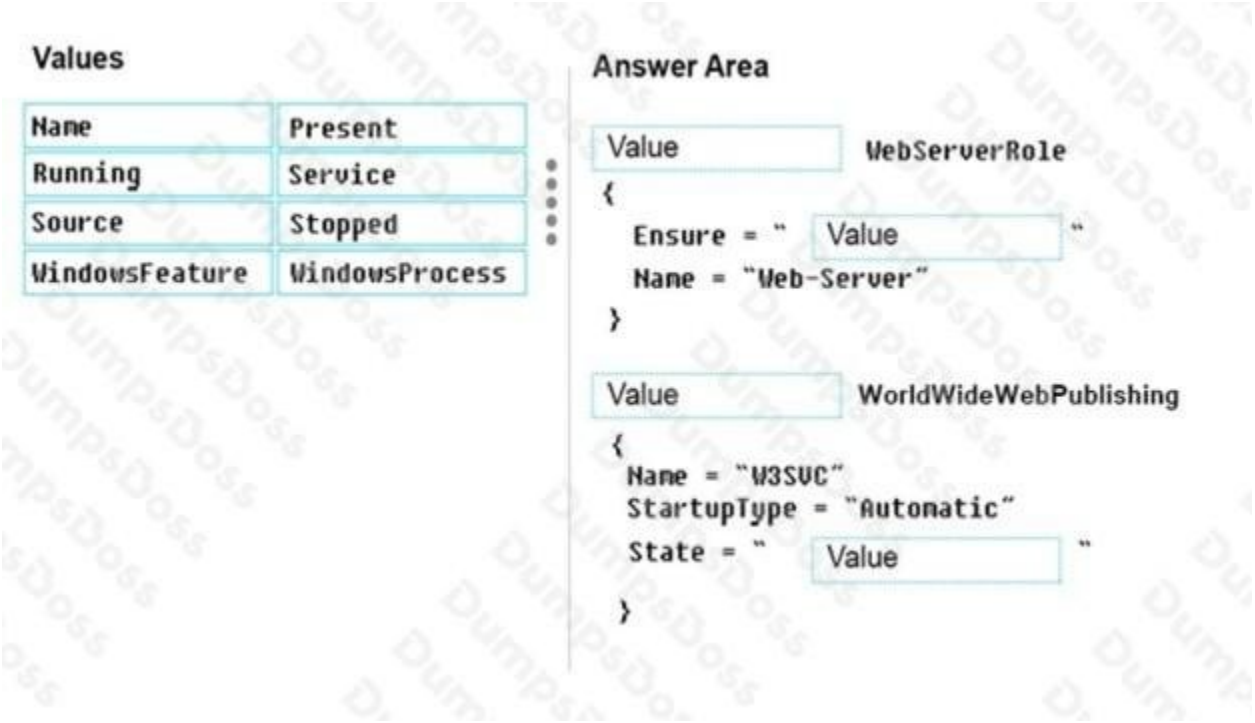
You have a network that contains several servers that run Windows Server 2016.

You need to use Desired State Configuration (DSC) to configure the servers to meet the following requirements: ▪ Install the Web Server server role

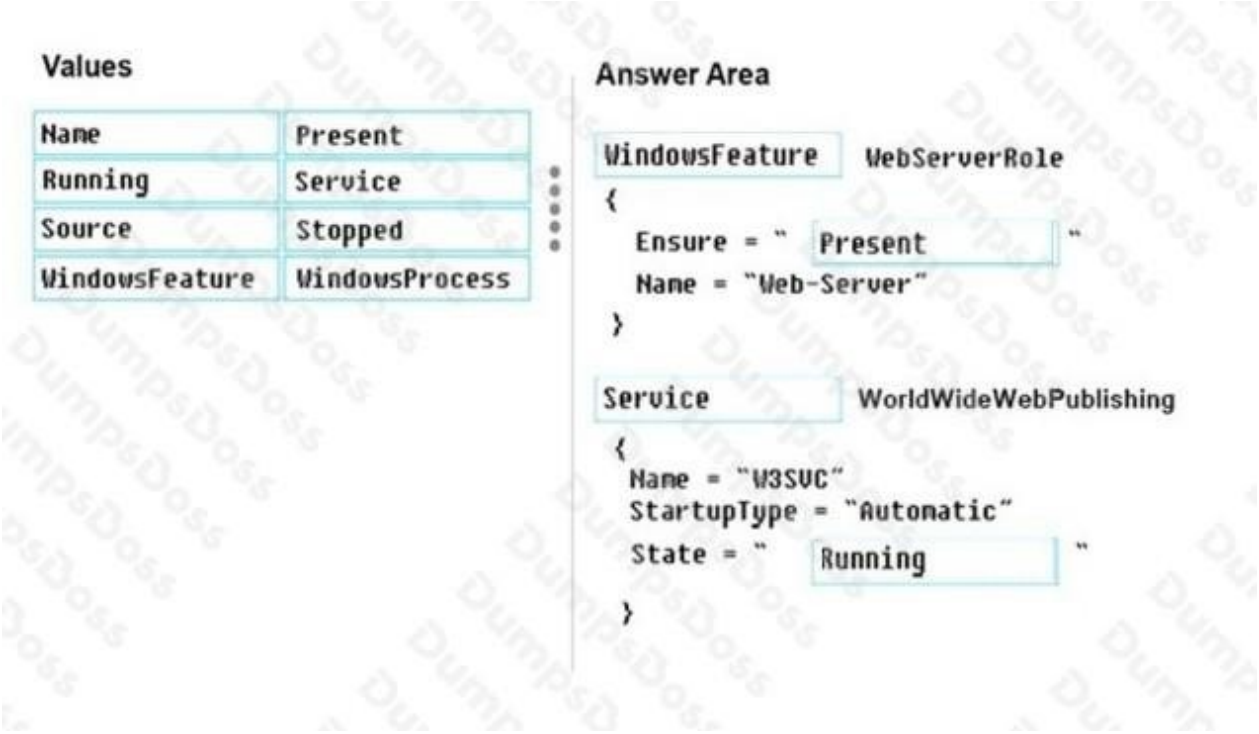
▪ Start the World Wide Web Publishing service

How should you configure the DSC recourses? To answer, drag the appropriate values to the correct locations. Each value 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.

Select and Place:



ANSWER:



Explanation:

Desired State Configuration uses resource blocks to define the state that a Windows Server node must maintain. For installation of the IIS Web Server role, the appropriate resource is `WindowsFeature`. This built-in DSC resource manages Windows Server roles and features by their feature name, and `Web-Server` is the feature name used for the Web Server role. Setting the `Ensure` property to `Present` declares that this role must be installed. DSC evaluates the node's current configuration and installs the role when it is not already available, allowing the configuration to remain idempotent on later runs.

The World Wide Web Publishing Service is managed separately because it is a Windows service rather than a server role. The `Service` DSC resource manages service configuration, including its service name, startup behavior, and operational state. The service name `W3SVC` identifies the World Wide Web Publishing Service. With `StartupType` set to `Automatic`, the service is configured to start when the server starts. Setting `State` to `Running` ensures that the service is actively started whenever DSC applies the configuration. Together, these resource declarations install the web-server capability and maintain the service needed to host IIS web content in a running state.

Microsoft's DSC documentation describes the `WindowsFeature` resource and its `Ensure` property at WindowsFeature DSC resource. The properties used to maintain Windows services, including `State` and `StartupType`, are documented at Service DSC resource.