PassQueen
a.k.a dumpsqueen.com

Configuring Windows Server Hybrid Advanced Services

Microsoft AZ-801

Version Demo

Total Demo Questions: 28

Total Premium Questions: 288

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Configure and manage Windows Server on-premises and hybrid infrastructure	121
Topic 2, Manage identity and security in Windows Server	78
Topic 3, Manage Windows Server networking	34
Topic 4, Manage Windows Server storage	36
Topic 5, Case Study 1	6
Topic 6, Case Study 2	4
Topic 7, Case Study 3	3
Topic 8, Case Study 4	4
Topic 9, Case Study 5	2
Total	288

QUESTION NO: 1

You have a file server named Server1 that runs Windows Server. Server1 hosts a shared folder named Finance.

You need to protect SMB traffic between client computers and Finance from disclosure and tampering.

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

NOTE: Each correct selection is worth one point.

A. Enable SMB encryption for the Finance share.

B. Require SMB signing on Server1.

C. Enable BranchCache on Server1.

D. Enable Data Deduplication on the Finance volume.

E. Configure Offline Files on the client computers.

ANSWER: A B

Explanation:

Configure **SMB encryption** for the Finance share to protect the confidentiality of data transferred through SMB. SMB encryption encrypts SMB data in transit, which prevents an attacker from reading the contents of files that are transferred between clients and the file server.

Configure **SMB signing** to protect the integrity of SMB messages. SMB signing adds a signature to SMB communications, allowing the receiving system to detect whether a message was modified in transit. Requiring both encryption and signing provides protection against disclosure and tampering for the SMB traffic.

For more information, see SMB security enhancements and Control SMB signing behavior.

QUESTION NO: 2 - (SIMULATION)

You have an on-premises server named Server1 that runs Windows Server and has the Hyper-V server role installed.

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

You plan to use Azure Backup.

Which type of vault should you configure for each resource? To answer, drag the appropriate vault types to the correct resources. Each vault 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: See the explanation for the answer

Explanation:

Configure a Recovery Services vault for both resources:

Server1 (on-premises Hyper-V): Recovery Services vault
VM1 (Azure virtual machine): Recovery Services vault

A Recovery Services vault supports Azure VM backups and on-premises Hyper-V workload protection through Azure Backup. A Backup vault is used for newer workload scenarios such as Azure Disk Backup and is not the vault type used to back up Hyper-V virtual machines or standard Azure VM backups.

QUESTION NO: 3

Your network contains an Active Directory Domain Services (AD DS) domain. The domain contains two servers named Server1 and Server2 that run Windows Server.

You need to ensure that you can use the Computer Management console to manage Server2. The solution must use the principle of least privilege.

Which two Windows Defender Firewall with Advanced Security rules should you enable on Server2? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. the COM+ Network Access (DCOM-In) rule

B. all the rules in the Remote Event Log Management group

C. the Windows Management Instrumentation (WMI-In) rule

D. the COM+ Remote Administration (DCOM-In) rule

E. the Windows Management Instrumentation (DCOM-In) rule

ANSWER: A B

Explanation:

Computer Management connects to a remote server through DCOM-based management components and uses Event Log RPC connectivity when the Event Viewer node is accessed. Enabling *the COM+ Network Access (DCOM-In) rule* permits the required inbound DCOM communication for remote COM-based administration without broadly opening unrelated remote-management services. Enabling *all the rules in the Remote Event Log Management group* permits the RPC, RPC endpoint mapper, and named-pipe communication required to read and manage the remote event logs through Computer Management.

This combination follows least privilege because it enables the specific firewall capabilities needed by Computer Management, including its Event Viewer functionality, rather than enabling a larger collection of general remote-administration rules. The Remote Event Log Management rules are intentionally grouped because remote event-log operations rely on more than one underlying inbound communication rule. Microsoft documents that remote server administration requires the applicable Windows Defender Firewall rule groups to be enabled on the managed server, and it documents the remote event-log firewall rule group for remote Event Viewer access. See Configure remote management in Server Manager and Configure Windows Defender Firewall.

QUESTION NO: 4

You have two servers named Server1 and Server2 that run Windows Server and have the Hyper-V server role installed.

You plan to use live migration to move virtual machines between Server1 and Server2.

Both servers have RDMA-capable network adapters.

You need to configure live migration to use SMB Direct.

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

NOTE: Each correct selection is worth one point.

A. Configure the live migration performance option to use SMB.

B. Enable RDMA on the adapters used for live migration.

C. Configure the live migration performance option to use compression.

D. Enable Hyper-V Replica on both servers.

E. Create a virtual Fibre Channel adapter for each virtual machine.

ANSWER: A B

Explanation:

You must configure the live migration performance option to use **SMB**. When SMB is selected for live migration, Hyper-V uses SMB 3.x for the memory-transfer traffic. SMB can use SMB Direct automatically when compatible RDMA network adapters are available.

You must also enable **RDMA** on the network adapters used for live migration. SMB Direct uses Remote Direct Memory Access to transfer data with reduced CPU overhead and high throughput. Both source and destination adapters must support and have RDMA enabled for SMB Direct to be used.

For more information, see Hyper-V live migration overview and SMB Direct.

QUESTION NO: 5

You have a server named Served that runs Windows Server. You install a custom app named App1 that is accessed by using TCP port 52310. Users report that they cannot access App1. You confirm that App1 is running on Server1.

You need to ensure that the users can access App1. The solution must only provide access to App1 on Server1. What should you do in Windows Defender Firewall with Advanced Security?

A. Create an inbound rule.

B. For the current profile, allow all inbound connections.

C. Create an outbound rule.

D. Create an isolation connection security rule.

ANSWER: A

Explanation:

Creating an inbound rule is the appropriate solution because App1 is listening for connections on TCP port 52310 on the server. Windows Defender Firewall evaluates unsolicited traffic arriving at the server as inbound traffic, and its default inbound behavior commonly blocks connections unless an allow rule applies. An inbound allow rule can be scoped specifically to TCP port 52310 and, if needed, to the applicable firewall profile, local IP address, remote IP addresses, or authorized users. This provides the required narrowly targeted access to the application hosted on Server1 without opening unrelated inbound services or ports. In Windows Defender Firewall with Advanced Security, create a new inbound rule of type Port, select TCP, specify local port 52310, choose Allow the connection, and apply it only to the profiles used by the server. Microsoft documents that inbound rules control traffic permitted to reach a computer, and that rules can be configured by protocol and local port. See Configure Windows Firewall and Windows Firewall rule settings.

QUESTION NO: 6

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

You plan to deploy a new line-of-business (LOB) application to VM1.

You need to ensure that the application can create child processes.

What should you configure on VM1?

A. Microsoft Defender Credential Guard

B. Microsoft Defender Application Control

C. Microsoft Defender SmartScreen

D. Exploit protection

ANSWER: D

Explanation:

Exploit protection is correct because it provides Windows process-mitigation settings that can be applied system-wide or to a specific executable. One of these mitigations is the child-process creation policy. This policy can prevent a protected application from launching child processes, which is useful for reducing exploitation techniques that rely on spawning another process. For the LOB application to create child processes, configure Exploit protection for the application executable and ensure that the child-process mitigation is not enabled for that application. This can be managed through Windows Security, Group Policy, imported Exploit protection configuration files, or the ProcessMitigations PowerShell module. A per-application configuration is generally preferable because it allows the required behavior for the LOB application without unnecessarily changing protections for unrelated workloads on VM1. Microsoft documents Exploit protection as the Windows feature used to configure and deploy these process mitigations: Enable exploit protection. The available configuration and management approach is also documented in the Set-ProcessMitigation PowerShell reference.

QUESTION NO: 7 - (HOTSPOT)

HOTSPOT

You have a Hyper-V failover cluster named Cluster1 at a main datacenter. Cluster1 contains two nodes that have the Hyper-V server role installed. Cluster1 hosts 10 highly available virtual machines.

You have a cluster named Cluster2 in a disaster recovery site. Cluster2 contains two nodes that have the Hyper-V server role installed.

You plan to use Hyper-V Replica to replicate the virtual machines from Cluster1 to Cluster2.

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

NOTE: Each correct selection is worth one point.

Hot Area:

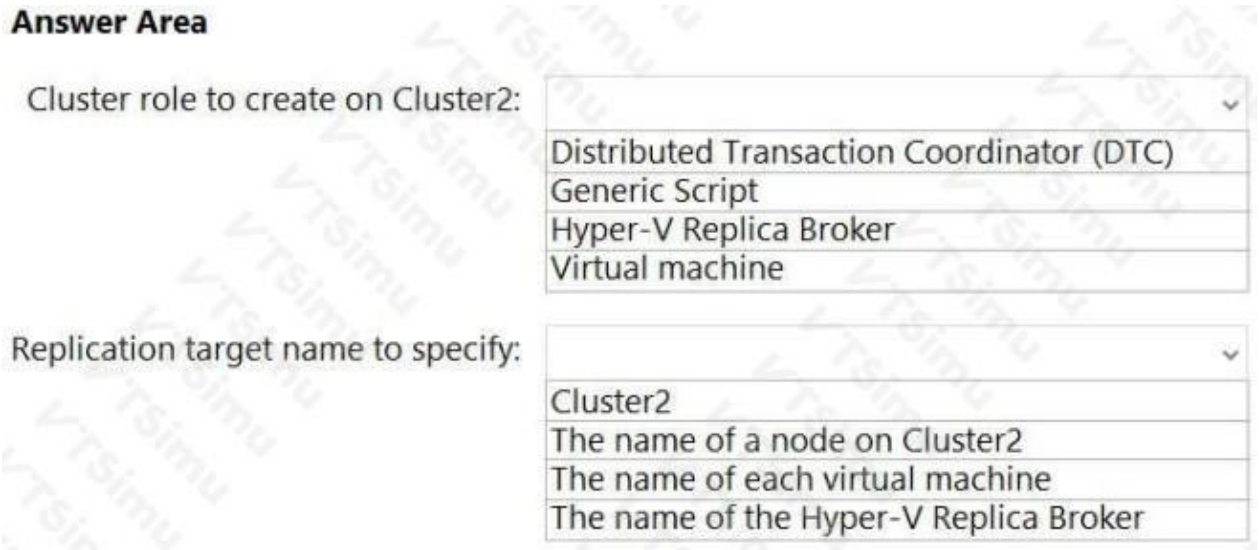

ANSWER:

Answer Area

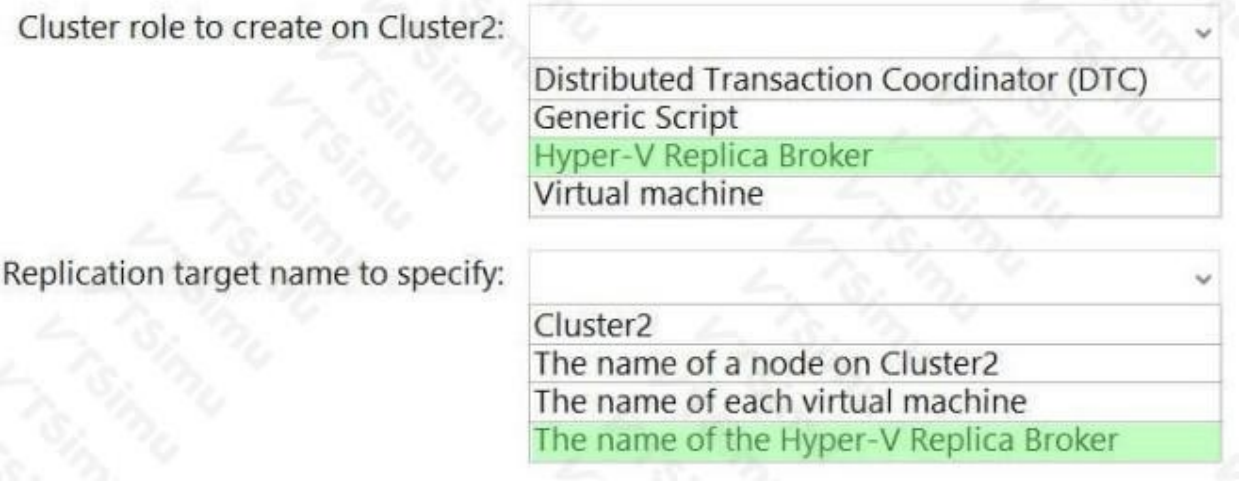

Explanation:

Cluster2 must be configured with the **Hyper-V Replica Broker** clustered role. This role provides the highly available Replica endpoint required when a Hyper-V failover cluster receives replicated virtual machines. Rather than tying replication to one physical Hyper-V host, the broker is a clustered service that can move between cluster nodes while retaining a stable client access point. This makes the disaster-recovery replication destination resilient to a node failure and allows Cluster2 to receive replicas as a cluster-based service.

After the broker role is created, the replication target must be specified by using **the name of the Hyper-V Replica Broker**. The broker is assigned a client access name and IP address as part of its clustered role configuration. That name is the logical network identity through which the source cluster communicates with the replica cluster. Using the broker name ensures that replication traffic continues to reach the destination service even when ownership of the broker role moves to the other node in Cluster2.

This configuration is the supported architecture for Hyper-V Replica where either the primary site, the replica site, or both use failover clusters. The broker centralizes the clustered Replica service and provides the consistent endpoint needed for highly available replica hosting. Microsoft's Hyper-V Replica documentation describes configuring the Replica Broker role and using its client access name for clustered replication scenarios. See Set up Hyper-V Replica and Failover clustering concepts for the clustered service and client-access-name model.

QUESTION NO: 8

You have two servers named Host1 and Host2 that run Windows Server 2022 and are members of a workgroup named Contoso. Both servers have the Hyper-V server role installed and are identical in hardware and configuration.

Host1 contains three virtual machines. Each server is located in a separate site that connect by using a high-speed WAN link.

You need to replicate the virtual machines from Host1 to Host2 for a disaster recovery solution.

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

NOTE: Each correct selection is worth one point.

A. Use certificate-based authentication and enable HTTPS exceptions.

B. Use Kerberos authentication and enable TCP exceptions.

C. Enable Hyper-V replication for each virtual machine.

D. Enable Host2 as a Replica server.

E. Enable live migration for Host1 and Host2.

F. Enable Host1 as a Replica server.

ANSWER: A C D

Explanation:

Because Host1 and Host2 are members of a workgroup rather than an Active Directory domain, Hyper-V Replica must use certificate-based authentication. Kerberos authentication depends on Active Directory and therefore is not appropriate for this workgroup deployment. A certificate that is trusted by both hosts provides the authentication required for replication traffic. Hyper-V Replica uses HTTPS when certificate-based authentication is selected, so the required HTTPS firewall exception must be enabled, normally for TCP port 443. Microsoft documents that certificate-based authentication is the supported authentication method for replica hosts that are not domain joined.

Host2 must be configured as a Replica server before it can receive replica virtual machines. Its Hyper-V Replica configuration defines that it accepts replication, uses certificate authentication, and stores incoming replica files in the selected location. Finally, replication must be enabled individually for each virtual machine on Host1. This configures Host1 as the primary host for those workloads, selects Host2 as the replica destination, and initiates the initial replication required to establish disaster-recovery copies. These steps provide asynchronous replication over the WAN without requiring a failover cluster or live migration configuration. See Set up Hyper-V Replica and Hyper-V Replica overview.

QUESTION NO: 9

You have a server named Server1 that runs Windows Server and has the Hyper-V server role installed. You have a Hyper-V failover cluster named Cluster1. All servers are members of the same domain.

You need to ensure that you use Hyper-V Replica with Kerberos authentication on the default port to replicate virtual machines from Cluster1 to Server1.

What should you do on Server1?

A. Add primary servers to the Hyper-V Replica Broker configuration.

B. From Hyper-V Settings, select Enable incoming and outgoing live migrations

C. From Windows Defender Firewall with Advanced Security, enable the Hyper-V Replica HTTPS Listener (TCP-In) rule.

D. From Windows Defender Firewall with Advanced Security, enable the Hyper-V Replica HTTP Listener (TCP-In) rule.

ANSWER: D

Explanation:

From Windows Defender Firewall with Advanced Security, enable the Hyper-V Replica HTTP Listener (TCP-In) rule is correct because Kerberos-based Hyper-V Replica communication uses HTTP rather than HTTPS. The default HTTP listener port for Hyper-V Replica is TCP port 80. Enabling this inbound firewall rule on Server1 permits replication traffic from the primary site, including the Hyper-V Replica Broker used by Cluster1, to reach Server1 on that default listener port.

Because every server is in the same Active Directory domain, Kerberos authentication can be used for mutual authentication between the primary and replica hosts. In the Hyper-V Replica configuration, Server1 must also be configured to accept replication as a replica server with Kerberos (HTTP) selected and the appropriate authorization entry for the cluster. The HTTP Listener firewall rule is the required network-access component that exposes the default Hyper-V Replica endpoint while retaining the intended Kerberos authentication model. Microsoft documents TCP 80 as the default port for Kerberos-authenticated Hyper-V Replica traffic and identifies the corresponding Hyper-V Replica HTTP listener firewall rule. See Set up Hyper-V Replica and Plan Hyper-V Replica.

QUESTION NO: 10

Your network contains an on-premises Active Directory Domain Services (AD DS) domain. The domain contains two virtual machines named VM1 and VM2 that run Windows Server.

You plan to implement a failover cluster named Cluster1 that will use VM1 and VM2 as nodes.

You need to ensure that Cluster1 can use floating IP addresses.

Which two components should you deploy? Each correct answer presents part-of the solution.

NOTE: Each correct selection is worth one point.

A. Network Load Balancing (NLB)

B. the Multipoint Services role

C. the Network Controller role

D. the Host Guardian Service role

E. Software Load Balancer (SLB)

ANSWER: C E

Explanation:

The Network Controller role and Software Load Balancer (SLB) should be deployed to provide floating IP functionality for workloads hosted on the failover cluster. In Windows Server Software-Defined Networking (SDN), SLB supplies virtual IP addresses that can represent a service independently of the individual virtual machines currently serving that workload. A floating IP can therefore remain reachable while the active clustered workload moves between VM1 and VM2. This enables clients to use a stable address rather than needing to identify the current active cluster node.

Network Controller is the centralized SDN management plane. It exposes the management interface used to configure and maintain SDN resources, including load-balancing policies and virtual IP configuration. SLB is the SDN data-plane component that implements the load-balancing and virtual-IP behavior on the network. Together, these components allow the environment to define the service address and apply the policies that direct traffic to the appropriate active backend instance. Microsoft documents SLB as the SDN component providing virtual IP-based load balancing, while Network Controller centrally manages SDN infrastructure and policies. See Software Load Balancer and Network Controller.

QUESTION NO: 11

You have a server named Server1 that runs Windows Server.

You need to ensure that only specific applications can modify the data in protected folders on Server1.

Solution: From App & browser control, you configure the Reputation-based protection.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No is correct because the required control is Microsoft Defender Controlled folder access, not Reputation-based protection. Controlled folder access protects designated folders from unauthorized changes by untrusted or unapproved applications. When it is enabled, applications must be trusted by Microsoft Defender Antivirus intelligence or explicitly added to the allowed-apps list before they can modify files in protected folders. This directly supports the requirement to permit only specific applications to write to protected locations on Server1.

Reputation-based protection is a different Microsoft Defender feature area. It uses reputation signals to help protect users from potentially unwanted apps, malicious or unrecognized applications, downloads, and websites; it does not define which applications may modify protected folders. To meet the stated requirement, configure Controlled folder access and, where necessary, allow the approved application executables using Microsoft Defender Antivirus policy, PowerShell, or the Windows Security interface. Microsoft documents Controlled folder access and its allowed-app mechanism at Controlled folder access. For the separate purpose and settings of reputation-based protection, see Protect your PC from potentially unwanted applications.

QUESTION NO: 12

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

VM1 is protected by Azure Backup.

You need to restore a single deleted file from the latest recovery point. The solution must minimize the time required to make the file available.

What should you use?

A. File recovery

B. Replace existing disks

C. Restore VM

D. Cross-region restore

ANSWER: A

Explanation:

File recovery is the appropriate Azure Backup restore option. File recovery lets you mount a recovery point from an Azure virtual machine backup and browse the protected disks to copy individual files and folders. This avoids restoring the entire virtual machine or replacing the existing disks.

After a recovery point is selected, Azure Backup provides a script that mounts the recovery point to a computer. You can locate the required file, copy it to the required destination, and then unmount the recovery point. This approach minimizes recovery time and avoids affecting the running VM.

For more information, see Restore files from an Azure virtual machine backup.

QUESTION NO: 13

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

You have a web server named Server1 that runs an IIS application by using the CONTOSO\WebAppSvc service account. The application connects to a Microsoft SQL Server instance that runs by using the CONTOSO\SqlSvc service account.

Users authenticate to the web application by using Kerberos authentication. You need to ensure that the application can access SQL Server by using the users' credentials.

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

NOTE: Each correct selection is worth one point.

A. Register an MSSQLSvc SPN for the SQL Server instance on CONTOSO\SqlSvc.

B. Configure CONTOSO\WebAppSvc for constrained delegation to the MSSQLSvc service.

C. Configure CONTOSO\SqlSvc for constrained delegation to the HTTP service.

D. Configure CONTOSO\WebAppSvc for unconstrained delegation.

ANSWER: A B

Explanation:

You must register the appropriate **MSSQLSvc service principal name (SPN)** on the SQL Server service account. The SPN enables clients and services to obtain Kerberos service tickets for the SQL Server service instance.

You must configure the web application's service account, **CONTOSO\WebAppSvc**, for Kerberos constrained delegation to the SQL Server MSSQLSvc service. This permits the IIS application to delegate a user's Kerberos credentials only to the specified back-end service.

Configuring delegation on the SQL Server service account does not allow the front-end web service to delegate credentials. Unconstrained delegation is not required and grants broader delegation capabilities than necessary.

For more information, see Setspn and Kerberos constrained delegation overview.

QUESTION NO: 14

You have a failover cluster named Cluster1 that hosts an application named App1.

The General tab in App1 Properties is shown in the General exhibit (Click the General Tat)

The failover tab in App1 Properties is shown in the failover exhibit (Click the Failover tab.)

Server 1 shuts down unexpectedly.

You need to ensure that when you start Server 1. App1 continues to run on Server2

Solution: From the General settings, you increase the priority of Server2 in the Preferred Owners list.

Does this meet the goal?

A. Yes

B. NO

ANSWER: B

Explanation:

NO is correct because Preferred Owners determines which cluster nodes are eligible to own the clustered role and establishes their ownership preference; it does not configure whether the role automatically moves back after the failed node returns. After Server1 unexpectedly shuts down, App1 can fail over to Server2. To ensure App1 remains on Server2 when Server1 starts again, the role's failback behavior must be configured to prevent failback, rather than relying on the Preferred Owners ordering on the General tab.

Failover Clustering exposes this control in the role's Failover properties. The failback policy determines whether a clustered role is allowed to move to its preferred node when that node rejoins the cluster, whether it is prevented from doing so, or whether it may do so during a defined time window. Selecting the setting that prevents failback keeps App1 online on Server2 following Server1's restart, unless an administrator deliberately moves the role or another event requires a failover. This separates node ownership preference from operational control over automatic post-recovery movement. See Microsoft's Failover Clustering documentation and Set-ClusterGroup reference.

QUESTION NO: 15

You have a server that runs Windows Server 2025 Standard and has the Hyper-V role installed. You need to upgrade the server to Windows Server 2025 Datacenter. The solution must minimize downtime. What should you run?

A. dism.exe

B. setup.exe

C. slagr.vbs

D. conwert.exe

ANSWER: A

Explanation:

dism.exe is the appropriate tool because Windows Server supports an online edition conversion from Standard to Datacenter by using Deployment Image Servicing and Management (DISM). For example, an administrator can use `DISM /Online /Set-Edition:ServerDatacenter /ProductKey:<product-key> /AcceptEula`, with a valid Windows Server 2025 Datacenter product key. This performs an edition conversion on the installed operating system rather than requiring a full operating-system reinstallation or server migration. The server must be restarted after the conversion is completed, but installed roles, including Hyper-V, and existing configuration remain in place. Therefore, the conversion confines service interruption principally to the required reboot and minimizes downtime for hosted workloads. Microsoft documents Standard-to-Datacenter conversion as a supported Windows Server edition-conversion path and identifies DISM as the command-line servicing tool used to set a new edition on an online image. Before performing the operation, ensure that maintenance is scheduled, virtual machines are shut down or live-migrated as appropriate, and a current backup is available. See Upgrade and conversion options for Windows Server and DISM Windows edition-servicing command-line options.

QUESTION NO: 16

Your network contains an Active Directory Domain Services (AD DS) domain. The domain contains domain controllers that run Windows Server 2019 and are configured as shown in the following table.

You plan to run the adprep /domainprep command.

Which domain controller should be available for the command to complete?

A. DC1

B. DC2

C. DC3

D. DC4

E. DC5

ANSWER: D

Explanation:

DC4 should be available because it holds the Infrastructure Master flexible single master operations (FSMO) role for the domain, as indicated by the configuration table. The `adprep /domainprep` operation performs domain-wide preparation by updating objects and permissions in the target domain. To make these changes safely and consistently, Adprep contacts the domain's Infrastructure Master role holder. Therefore, the command requires connectivity to, and availability of, DC4 for successful completion.

Domain preparation is normally performed after forest preparation when extending Active Directory for a newer Windows Server version or for software that requires directory-schema and domain-level updates. The account running the operation must also have the required permissions, typically membership in the Domain Admins group for the target domain. Although directory replication distributes the resulting updates to other domain controllers after the operation completes, the Infrastructure Master must be operational at the time `adprep /domainprep` runs.

Microsoft documents the Adprep command and its domain-preparation behavior at Adprep command reference. For additional background on FSMO role placement and the Infrastructure Master role, see Operations master roles.

QUESTION NO: 17 - (HOTSPOT)

Your network contains an Active Directory Domain Services (AD DS) domain The domain contains an organizational unit (OU) named OU1 and a user named User1.

You plan to deploy a Hyper V failover cluster named Cluster1.

You need to prestage the account for Cluster! and ensure that User1 can deploy Cluster1. The solution must follow the principle of least privilege.

Which action should you perform, and which permissions should you grant to User1 for Cluster1? To answer, select the appropriate options in the answer area

NOTE: Each correct selection is worth one point.

ANSWER:

Explanation:

A Hyper-V failover cluster uses a computer object in Active Directory Domain Services called the Cluster Name Object (CNO). The CNO has the same name as the cluster, so the account to prestage must be a **computer account named Cluster1**. Pre-staging is performed by creating that computer object in the appropriate OU and then disabling it before the cluster is created. During deployment, the failover clustering process uses the prestaged object rather than attempting to create a new CNO.

User1 must be able to take control of and configure that specific prestaged computer object as part of creating Cluster1. Granting **Full control** to User1 on the Cluster1 computer object supplies the rights needed to enable the object and manage the required directory attributes during cluster deployment. Because the permission is applied to the individual Cluster1 account rather than broadly at the OU or domain level, the access is constrained to the cluster object being deployed and follows least-privilege administration.

Microsoft documents that a prestaged CNO is created as a disabled computer account and that the account used to create the cluster requires Full Control permissions for that object. See Prestage cluster computer objects in Active Directory Domain Services and Create a failover cluster.

QUESTION NO: 18 - (DRAG DROP)

You need to meet the technical requirements f or Cluster2.

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

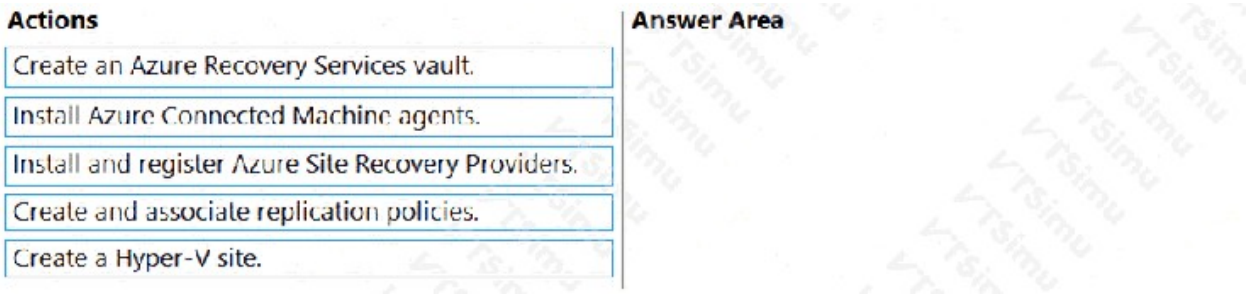

ANSWER:

Explanation:

To prepare Cluster2 for Azure Site Recovery replication, first create an Azure Recovery Services vault. The vault is the Azure management resource that holds the Site Recovery configuration, protected-item information, recovery points, and replication settings. It must exist before the on-premises Hyper-V infrastructure can be registered for protection.

Next, create a Hyper-V site in the vault. A site provides the logical representation of the Hyper-V environment that will participate in replication. This lets Azure Site Recovery organize the hosts in Cluster2 and apply consistent Site Recovery configuration to that environment.

Then install and register the Azure Site Recovery Provider on the relevant Hyper-V hosts. During registration, each host connects to the Recovery Services vault and is associated with the configured Hyper-V site. This registration establishes the trusted management connection that Azure Site Recovery uses to discover and protect workloads on the hosts.

Finally, create and associate replication policies. A replication policy defines the operational replication behavior, including settings such as recovery-point retention and application-consistent snapshot frequency. Associating the policy with the Hyper-V site makes those protection settings available to the registered hosts, completing the required configuration before replication is enabled for virtual machines. This sequence follows Microsoft's Hyper-V-to-Azure Site Recovery setup flow. For implementation details, see Set up disaster recovery of on-premises Hyper-V VMs to Azure and Prepare Azure for disaster recovery of on-premises Hyper-V VMs.

QUESTION NO: 19

Your network contains an Active Directory Domain Services (AD DS) domain that has domain controllers in multiple sites.

You discover that changes made to group membership on DC1 do not appear on DC3.

You need to identify the replication partners for DC3 and determine whether replication failures occur elsewhere in the domain.

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

NOTE: Each correct selection is worth one point.

A. repadmin /showrepl DC3

B. repadmin /replsummary

C. repadmin /syncall DC3

D. dcdiag /test:DNS

ANSWER: A B

Explanation:

repadmin /showrepl DC3 displays the inbound replication partners for DC3 and the result of the most recent replication attempt from each partner. This helps identify which connection is failing and provides the associated replication error.

repadmin /replsummary summarizes replication failures across domain controllers. It provides a domain-wide view that can help determine whether the issue is isolated to DC3 or affects other replication paths.

The `repadmin /syncall` command initiates replication rather than reporting the existing replication state. The DNS test in Dcdiag can help investigate DNS issues, but it does not provide the required replication partner and failure summary.

For more information, see Repadmin /showrepl and Repadmin /replsummary.

QUESTION NO: 20 - (DRAG DROP)

You have an on-premises server named Server1 that runs Windows Server.

You have an Azure subscription that uses Microsoft Defender for Cloud.

You need to onboard Server1 to Defender for Cloud.

Which actions should you perform in sequence? To answer, drag the appropriate actions to the correct order. Each action 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.

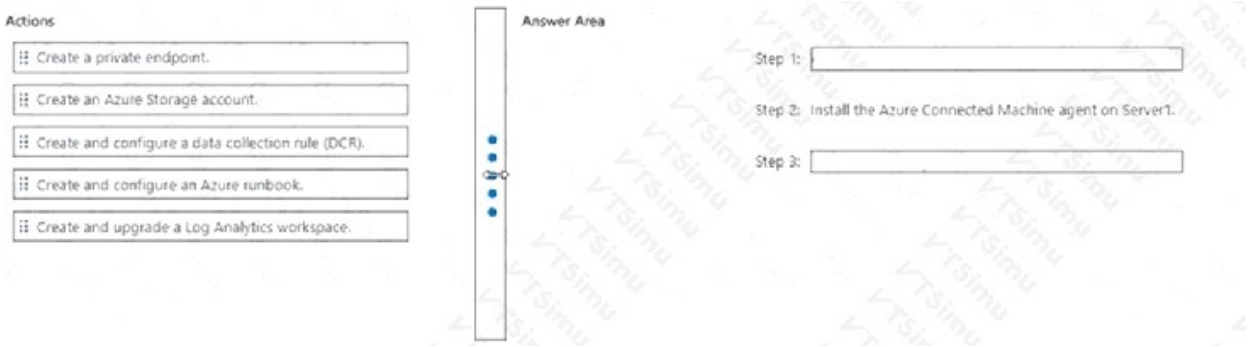

ANSWER:

Explanation:

Start by creating and upgrading a Log Analytics workspace. The workspace provides the Azure-based destination for collected monitoring and security-related log data. It is important to have this destination established before configuring collection, because the data collection configuration needs a workspace to which it can send data. A workspace also gives Defender for Cloud and associated Azure monitoring services a central location in which to store, query, and analyze the information collected from the server.

Next, install the Azure Connected Machine agent on Server1. This onboards the on-premises Windows Server machine to Azure Arc, representing the server as an Azure resource. Once connected, Server1 can be centrally managed and can receive supported Azure extensions and monitoring configurations. Microsoft documents the Connected Machine agent as the component used to connect physical servers and virtual machines outside Azure to Azure Arc. See Azure Connected Machine agent overview.

Finally, create and configure a data collection rule (DCR). A DCR specifies what monitoring data is collected from the connected server and identifies the Log Analytics workspace destination. This completes the collection configuration after both the destination workspace and the Azure-connected server are available. Microsoft's data collection rule overview explains that DCRs define the data collected by Azure Monitor and the destinations to which that data is sent. With Server1 connected through Azure Arc and its collection data routed to the workspace, Defender for Cloud can use the server's Azure-connected monitoring and security configuration as part of protecting the hybrid environment.

QUESTION NO: 21

You have 50 on-premises servers that run Windows Server. The servers are enabled for Azure Arc and have the Azure Monitor agent installed.

You need to collect System event logs and processor performance data from the servers to a Log Analytics workspace.

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

NOTE: Each correct selection is worth one point.

A. Create a data collection rule that includes Windows event logs and performance counters.

B. Associate the data collection rule with the Azure Arc-enabled servers.

C. Create an Azure Automation account.

D. Install the Log Analytics agent on each server.

ANSWER: A B

Explanation:

A **data collection rule (DCR)** defines the data sources, such as Windows events and performance counters, and specifies the destination Log Analytics workspace. The DCR must contain the required event-log and performance-counter collection settings.

You must also **associate the DCR with the Azure Arc-enabled servers**. The Azure Monitor agent uses the association to retrieve the collection configuration and send the selected data to the configured destination.

Creating a Log Analytics workspace alone does not define what data the agent collects. The legacy Log Analytics agent is not required when the Azure Monitor agent is already installed.

For more information, see Data collection rules in Azure Monitor.

QUESTION NO: 22 - (SIMULATION)

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

You need to perform the following tasks:

• Createasnapshotofcontoso.com.

• Expose the snapshot as a read-only AD DS instance.

What should you use for each task? To answer, drag the appropriate tools to the correct tasks. Each tool 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: See the explanation for the answer

Explanation:

Correct tool assignments:

Create a snapshot of contoso.com: `ntdsutil.exe`
Expose the snapshot as a read-only AD DS instance: `dsamain.exe`

Use `ntdsutil` to create and mount/manage an AD DS database snapshot. Then use `dsamain` against the mounted snapshot database (`ntds.dit`) to publish it on an unused LDAP port as a read-only directory service instance.

```
ntdsutil
activate instance ntds
snapshot
create

:: After mounting the snapshot, expose its database:
dsamain -dbpath <mounted_snapshot_path>\Windows\NTDS\ntds.dit -ldapport 51389
```

QUESTION NO: 23

Your on-premises network contains an Active Directory Domain Services (AD DS) domain. The domain contains a file server named Server 1 that hosts multiple shared folders.

You have an Azure subscription.

You need to migrate the files stored on Server! to Azure by using an Azure Data Box gateway.

To which storage services can you migrate the files?

A. Azure Files only

B. Azure files and Azure Storage page blobs only

C. Azure Files and Azure Storage block blobs only

D. Azure Files, Azure Storage page blobs, and Azure Storage block blobs

ANSWER: D

Explanation:

Azure Files, Azure Storage page blobs, and Azure Storage block blobs is correct because Azure Data Box Gateway is a virtual appliance that provides a local SMB or NFS endpoint for on-premises workloads and transfers the ingested data to Azure storage. The gateway supports cloud storage accounts configured for Azure Blob Storage or Azure Files. For Blob Storage, it supports both block blob and page blob destinations. This enables a file server migration to use Azure Files where a managed SMB file-share service is required, or Blob Storage where data is to be stored as block blobs or page blobs. The appropriate destination is selected when the share is created on the Data Box Gateway, which maps the local share to the relevant Azure storage resource. Microsoft's Data Box Gateway documentation describes its role in transferring data to Azure and its support for Azure Blob Storage and Azure Files. Microsoft's Blob Storage documentation also identifies block blobs and page blobs as supported blob types. See Azure Data Box Gateway overview and Introduction to Azure Blob Storage.

QUESTION NO: 24

Your network contains an Active Directory Domain Services (AD DS) domain. The functional level of the domain is Windows Server 2012 R2.

You need to create and link an authentication policy silo.

What should you do first?

A. Raise the domain functional level to Windows Server 2016.

B. Raise the domain functional level to Windows Server 2025.

C. Create an authentication policy.

D. Create an authentication policy silo.

ANSWER: C

Explanation:

Create an authentication policy first. Authentication policies and authentication policy silos are Active Directory features introduced with Windows Server 2012 R2, so a domain functional level of Windows Server 2012 R2 already meets the prerequisite; no functional-level increase is required. An authentication policy defines the sign-in and credential restrictions that will apply to user, computer, or service accounts. The authentication policy silo then provides the boundary used to group those accounts and associate the appropriate authentication policies with them. After the relevant policy or policies exist, an administrator can create the silo, link the policies to the silo as appropriate, add the accounts that should be governed by it, and enforce the policy when ready. This sequencing ensures that the silo has an authentication-policy configuration available to apply to its members. Microsoft documents authentication policies and silos as the mechanism for

restricting how privileged accounts authenticate and describes their Windows Server 2012 R2 availability and configuration model. See Authentication Policies and Authentication Policy Silos and New-ADAuthenticationPolicy.

QUESTION NO: 25

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

You create an Azure Storage Sync Service and register Server1 and Server2.

You need to synchronize the contents of D:\Data on both servers with an Azure file share.

Which two endpoint types should you create in the sync group? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. a cloud endpoint

B. a server endpoint

C. an Azure file share snapshot

D. a DFS Namespace target

ANSWER: A B

Explanation:

A **cloud endpoint** links an Azure file share to an Azure File Sync sync group. Each sync group has one cloud endpoint, which provides the central Azure file share used by the synchronization relationship.

A **server endpoint** links a path on a registered Windows Server, such as D:\Data, to the sync group. Create a server endpoint for the applicable path on each server that must synchronize with the cloud endpoint.

An Azure file share snapshot is a point-in-time copy for recovery and does not establish synchronization. A DFS Namespace target provides a namespace referral target and is unrelated to Azure File Sync endpoint configuration.

For more information, see Deploy Azure File Sync.

QUESTION NO: 26

You have a server named Server1 that runs Windows Server.

You need to configure Windows Defender Firewall logging on Server1. The logs must contain records of dropped packets and successful connections.

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

NOTE: Each correct selection is worth one point.

A. Log dropped packets

B. Log successful connections

C. Display a notification when Windows Defender Firewall blocks a new app

D. Require authentication for inbound and outbound connections

ANSWER: A B

Explanation:

Enable **Log dropped packets** to record packets that Windows Defender Firewall blocks. This information is useful when troubleshooting firewall rules or identifying unexpected connection attempts.

Enable **Log successful connections** to record connections that the firewall permits. Together, these settings provide visibility into both blocked and allowed network activity for the applicable firewall profile.

Firewall notifications display prompts to an interactive user and do not provide the required event logging. Connection security rules define IPsec behavior and do not enable firewall packet logging.

For more information, see Configure Windows Defender Firewall.

QUESTION NO: 27

Your company uses Storage Spaces Direct.

You need to view the available storage in a Storage Space Direct storage pool.

What should you use?

A. System Configuration

B. File Server Resource Manager (FSRM)

C. the Get-StorageFileServer cmdlet

D. Failover Cluster Manager

ANSWER: D

Explanation:

Failover Cluster Manager is the appropriate management console for viewing capacity information in a Storage Spaces Direct deployment. Storage Spaces Direct is integrated with Windows Server Failover Clustering, and Failover Cluster Manager exposes the cluster's storage resources in a single administrative view. After connecting to the cluster, expand *Storage* and select *Pools* to view the Storage Spaces Direct pool. The details pane displays pool capacity information, including allocated or used capacity and free capacity, allowing an administrator to determine the available storage in the pool.

This console is particularly useful because it presents the storage pool in its cluster context, alongside related virtual disks, physical disks, and cluster roles. It provides an at-a-glance operational view without requiring a separate reporting configuration or a custom PowerShell query. Microsoft identifies Failover Cluster Manager as a core tool for administering failover clusters, including their clustered storage resources. Storage Spaces Direct uses a pool of local server drives to provide resilient, software-defined storage across the cluster, so pool-level capacity is the relevant measure for this requirement.

For more information, see Failover Clustering overview and Storage Spaces Direct overview.

QUESTION NO: 28

Your network contains an Active Directory Domain Services (AD DS) domain that has a Windows Server 2012 R2 forest and domain functional level. The domain contains the domain controllers shown in the following table.

Name	Operating system
DC1	Windows Server 2012 R2
DC2	Windows Server 2016
DC3	Windows Server 2016

You need to ensure that you can deploy a new domain controller that runs Windows Server 2025. The solution must minimize administrative effort.

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

NOTE: Each correct selection is worth one point.

A. Upgrade DC1, DC2, and CO to Windows Server 2025.

B. Raise the forest and domain functional level to Windows Server 2025.

C. Raise the forest and domain functional level to Windows Server 2016.

D. Remove DC2 and DC3.

E. Remove DC 1.

ANSWER: C E

Explanation:

To introduce a Windows Server 2025 domain controller, the Active Directory forest and domain must use at least the Windows Server 2016 functional level. Windows Server 2025 does not support use in a forest or domain that remains at the Windows Server 2012 R2 functional level. Therefore, raising both functional levels to Windows Server 2016 establishes the required AD DS compatibility baseline while avoiding an unnecessary operating-system upgrade of every existing domain controller.

Before the functional levels can be raised, all domain controllers that run an operating system earlier than Windows Server 2016 must be removed from the domain. Removing DC 1 removes the domain controller that prevents raising the functional levels. After ensuring that DC 1 holds no required roles and has been appropriately demoted, the domain functional level can be raised first, followed by the forest functional level. The environment can then support deployment of the Windows Server 2025 domain controller.

Microsoft documents the supported Active Directory functional level requirements and the process for raising them in Active Directory Domain Services functional levels and Raise the domain and forest functional levels.