PassQueen
a.k.a dumpsqueen.com

Planning and Administering Microsoft Azure for SAP Workloads

Microsoft AZ-120

Version Demo

Total Demo Questions: 35

Total Premium Questions: 357

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Design and implement infrastructure to support SAP workloads on Azure	121
Topic 2, Design and implement SAP workloads on Azure	23
Topic 3, Migrate SAP workloads to Azure	73
Topic 4, Maintain SAP workloads on Azure	75
Topic 5, Design and implement high availability and disaster recovery	51
Topic 6, Mix Questions	3
Topic 7, Case Study 1	3
Topic 8, Case Study 2	3
Topic 9, Case Study 3	3
Topic 10, Case Study 4	2
Total	357

QUESTION NO: 1

You deploy an SAP production landscape on Azure virtual machines that run SUSE Linux Enterprise Server (SLES).

You need to generate a report that details performance differences between instances of an SAP AS ABAP system.

What should you use?

A. JMeter

B. Micro Focus LoadRunner

C. SAP UI5 SpeedtestTool

D. ABAPmeter

ANSWER: D

Explanation:

ABAPMETER is the appropriate tool for comparing performance across instances in an SAP NetWeaver AS ABAP system. It is an SAP-delivered ABAP report designed to execute a standardized set of operations on the selected application-server instances and present comparable runtime measurements. This makes it useful for identifying instance-level performance variation, such as differences caused by compute capacity, configuration, resource contention, or network-related conditions in an Azure-hosted SAP landscape. The report provides a controlled comparison because the same workload is run against each instance, allowing administrators to assess the relative health and responsiveness of the ABAP application servers rather than relying only on operating-system-level metrics. For Azure SAP workload investigations, ABAPMETER is therefore a relevant application-layer diagnostic tool and can complement Azure monitoring data when isolating where a performance discrepancy occurs. Microsoft's SAP workload performance guidance discusses using SAP-native performance analysis tools for application-level troubleshooting, while SAP documentation specifically identifies ABAPMETER as a mechanism for measuring and comparing application-server performance. See Microsoft Learn: SAP application performance standard scenarios and SAP KBA 2879613.

QUESTION NO: 2

You are migrating SAP to Azure. The ASCS application servers are in one Azure zone, and the SAP database

server in in a different Azure zone. ASCS/ERS is configured for high availability.

During performance testing, you discover increased response times in Azure, even though the Azure

environment has better computer and memory configurations than the on-premises environment.

During the initial analysis, you discover an increased wait time for Enqueue.

What are three possible causes of the increased wait time? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. a missing Enqueue profile

B. disk I/O during Enqueue backup operations

C. misconfigured load balancer rules and health check probes for Enqueue and ASCS

D. active Enqueue replication

E. network latency between the database server and the SAP application servers

ANSWER: C D E

Explanation:

Misconfigured load balancer rules and health check probes for Enqueue and ASCS can affect reachability of the highly available ASCS virtual IP and can introduce connection disruption or delays while SAP application servers access central services. The load-balancer configuration is therefore an essential part of a functioning ASCS/ERS high-availability design.

Active Enqueue replication can also contribute to response-time overhead because ASCS/ERS maintains replication of enqueue state to support failover. This traffic and processing must be considered when investigating enqueue-related latency, particularly under a high lock-request workload.

Network latency between the database server and the SAP application servers is especially relevant in a zonal deployment. Azure availability zones are physically separate locations, and inter-zone traffic introduces latency. SAP application processing is sensitive to latency between application and database tiers; slower database round trips can extend transaction processing and the time during which application work waits for or holds enqueue resources. Microsoft specifically advises validating whether inter-zone latency is acceptable for the SAP workload in the selected region. Review the ASCS/ERS load-balancer configuration and zonal latency characteristics as part of performance validation. See SAP high availability with availability zones and SAP high-availability guidance for Azure.

QUESTION NO: 3

You have an existing SAP production landscape on Azure. The SAP application virtual machines have static IP addresses.

You need to replicate the virtual machines to another Azure region by using Azure Site Recovery. The source and target subnets have different address ranges.

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

NOTE: Each correct selection is worth one point.

A. Stop the virtual machines.

B. Create a backup policy.

C. Create a recovery plan.

D. Modify the networking configuration.

E. Replicate the virtual machines.

ANSWER: C D E

Explanation:

“Modify the networking configuration,” “Create a recovery plan,” and “Replicate the virtual machines” are required parts of an Azure Site Recovery disaster-recovery design for SAP application VMs when the recovery-region subnet uses a different address range. Before enabling replication, configure the target networking details so that recovered VMs connect to the appropriate target virtual network and subnet. Because the source static private IP addresses cannot be used unchanged in a subnet with a different address prefix, define the appropriate target-side addressing and network configuration for the recovered workload.

A recovery plan provides the coordinated failover workflow needed for a multi-VM SAP landscape. It lets administrators group workload components, control the order in which they start, and use automation steps where needed to apply post-failover configuration such as SAP- or network-dependent updates. This supports a consistent, repeatable recovery process rather than failing over application servers independently.

Finally, replication must be enabled for the virtual machines so that Azure Site Recovery continuously copies the protected VM data to the target region and makes test failover, planned failover, and disaster recovery possible. Microsoft documents Azure-to-Azure replication configuration at Enable replication for Azure VMs and recovery-plan orchestration at Recovery plan overview.

QUESTION NO: 4

You have an on-premises datacenter.

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

VNet1 uses forced tunneling to force all outbound traffic to the internet via the datacenter.

You deploy an SAP landscape to VNet1. The landscape uses virtual machines that run Red Hat Enterprise Linux (RHEL) and uses the pay-as-you-go rates.

You discover that the virtual machines fail to receive operating system updates.

You need to ensure that the virtual machines receive operating system updates.

What should you do?

A. Add user-defined routes (UDRs) to the route table.

B. Deploy Azure Update Manager.

C. Deploy an Azure firewall.

D. Configure an Azure DNS Private Resolver.

ANSWER: A

Explanation:

Add user-defined routes (UDRs) to the route table. Pay-as-you-go RHEL virtual machines in Azure obtain entitled operating-system packages through Red Hat Update Infrastructure (RHUI). Forced tunneling replaces the normal default outbound path with a route through the on-premises network. Therefore, the route table must include the required UDR exception(s) so RHUI update traffic can reach the appropriate Azure-hosted RHUI endpoints rather than being sent along an unusable forced-tunnel path. The precise destination prefixes and next hop should be configured according to the RHUI guidance and the organization's network design; the key requirement is to preserve a valid outbound path for the RHEL repository service while retaining forced tunneling for other applicable traffic. Once the route exception is associated with the SAP VM subnet, the RHEL package manager can contact RHUI and retrieve supported security and operating-system updates. Microsoft documents that on-demand RHEL images use RHUI for updates and describes networking considerations for those images in the Red Hat Update Infrastructure documentation. For the route-selection behavior that makes this solution effective, see Azure user-defined routes documentation.

QUESTION NO: 5

You plan to deploy SAP HANA on an Azure virtual machine that supports accelerated networking.

You need to minimize network latency and CPU utilization for network traffic between the SAP HANA database server and the SAP application servers.

What should you do?

A. Enable accelerated networking on the network interfaces.

B. Enable IP forwarding on the network interfaces.

C. Deploy an Azure Application Gateway.

D. Configure a Basic SKU public IP address.

ANSWER: A

Explanation:

Enable **accelerated networking** on the virtual machine network interfaces. Accelerated networking uses single root I/O virtualization (SR-IOV) to bypass the virtual switch for supported network traffic. This lowers latency, reduces jitter, and decreases CPU utilization for network processing.

For SAP workloads, this is beneficial for the latency-sensitive communication between SAP application servers and SAP HANA. Accelerated networking must be enabled on a supported VM size and operating system image. The VM is normally stopped and deallocated before the setting is enabled on an existing network interface.

For more information, see What is accelerated networking? and Network configuration for SAP workloads.

QUESTION NO: 6

You are planning the Azure network infrastructure to support the disaster recovery requirements.

What is the minimum number of virtual networks required for the SAP deployment?

A. 1

B. 2

C. 3

D. 4

ANSWER: B

Explanation:

2 is correct because regional disaster recovery requires a production environment in one Azure region and a recoverable target environment in a separate Azure region. An Azure virtual network is scoped to a single region, so the network that hosts the production SAP workload cannot also serve as the virtual network in the disaster-recovery region. Consequently, the minimum design uses one virtual network for the primary SAP deployment and one virtual network for the secondary, recovery-region deployment.

Azure Site Recovery supports Azure-to-Azure replication, including selecting a target region and target virtual network for replicated machines. During a regional outage, the replicated SAP application and database servers can be failed over into the secondary region and connected to the virtual network deployed there. This separation provides the regional fault-domain isolation required by the stated objective that production databases withstand an Azure-region failure. Network connectivity between the primary and recovery environments, such as global virtual network peering or VPN/ExpressRoute designs, can be added as required, but it does not eliminate the need for a virtual network in each region. See Azure-to-Azure disaster recovery with Site Recovery and Azure Virtual Network FAQ.

QUESTION NO: 7

You deploy several SAP application servers on Azure virtual machines.

An application running on each server must retrieve the same database credential. You must prevent the credential from being stored on the virtual machines and minimize credential-management overhead.

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

NOTE: Each correct selection is worth one point.

A. Store the credential as a secret in Azure Key Vault.

B. Create a user-assigned managed identity and assign it to the virtual machines.

C. Store the credential in an Azure Automation variable.

D. Create a local administrator account that has the same password on every virtual machine.

E. Store the credential in a network security group rule.

ANSWER: A B

Explanation:

Store the database credential as a **secret in Azure Key Vault**. Key Vault is designed to protect secrets, keys, and certificates, and it provides centralized management and auditing for sensitive application credentials.

Create a **user-assigned managed identity** and assign it to each SAP application-server virtual machine. The same identity can be granted only the required permission to retrieve the secret from Key Vault. The application can then authenticate to Key Vault without storing a password, certificate, or service-principal secret on the virtual machines.

For more information, see Azure Key Vault overview and Managed identities for Azure resources.

QUESTION NO: 8

You plan to deploy SAP HANA on Azure virtual machines.

You need to verify that a proposed virtual machine series is certified and supported by SAP for the required SAP HANA deployment.

What should you review?

A. the SAP HANA certified IaaS platforms list

B. the Azure Advisor cost recommendations

C. the Azure virtual machine SLA page

D. the SAP Quick Sizer report

ANSWER: A

Explanation:

You should review the **SAP HANA certified IaaS platforms list**. SAP publishes the supported cloud infrastructure configurations for SAP HANA, including the certified Azure virtual machine types, supported operating systems, and applicable SAP HANA deployment scenarios.

An Azure virtual machine size can provide sufficient vCPU and memory capacity but still be unsuitable if it is not certified for the intended SAP HANA configuration. The supported configuration must be validated against SAP certification information and the relevant Microsoft SAP HANA guidance before deployment.

For more information, see SAP HANA certified IaaS platforms and SAP HANA operations on Azure virtual machines.

QUESTION NO: 9

You deploy an SAP HANA scale-out system on Azure virtual machines.

You need to provide shared storage for the `/hana/shared` file system. The storage solution must support NFS and provide high throughput for the SAP HANA nodes.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure NetApp Files capacity pool and volume.

B. Configure the shared volume to use NFSv4.1.

C. Use a locally attached temporary disk.

D. Enable read caching on the operating system disk.

E. Use an Azure Files SMB share.

ANSWER: A B

Explanation:

Create an **Azure NetApp Files capacity pool** and a volume for the shared SAP HANA file system. Azure NetApp Files provides high-performance NFS storage and is supported for SAP HANA shared file systems in applicable SAP architectures.

You should configure the volume to use **NFSv4.1** when the SAP HANA deployment and operating system support that protocol. The HANA nodes can then mount the same Azure NetApp Files volume as `/hana/shared`.

Azure NetApp Files capacity and service level must be sized to meet the throughput and latency requirements of the workload. For more information, see SAP HANA scale-out with standby node on Azure NetApp Files and What is Azure NetApp Files?.

QUESTION NO: 10 - (HOTSPOT)

HOTSPOT

You have an on-premises SAP environment.

Backups are performed by using tape backups. There are 50 TB of backups.

A Windows file server has BMP images of checks used by SAP Finance. There are 9 TB of images.

You need to recommend a method to migrate the images and the tape backups to Azure. The solution must maintain continuous replication of the images.

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

NOTE: Each correct selection is worth one point.

Hot Area:

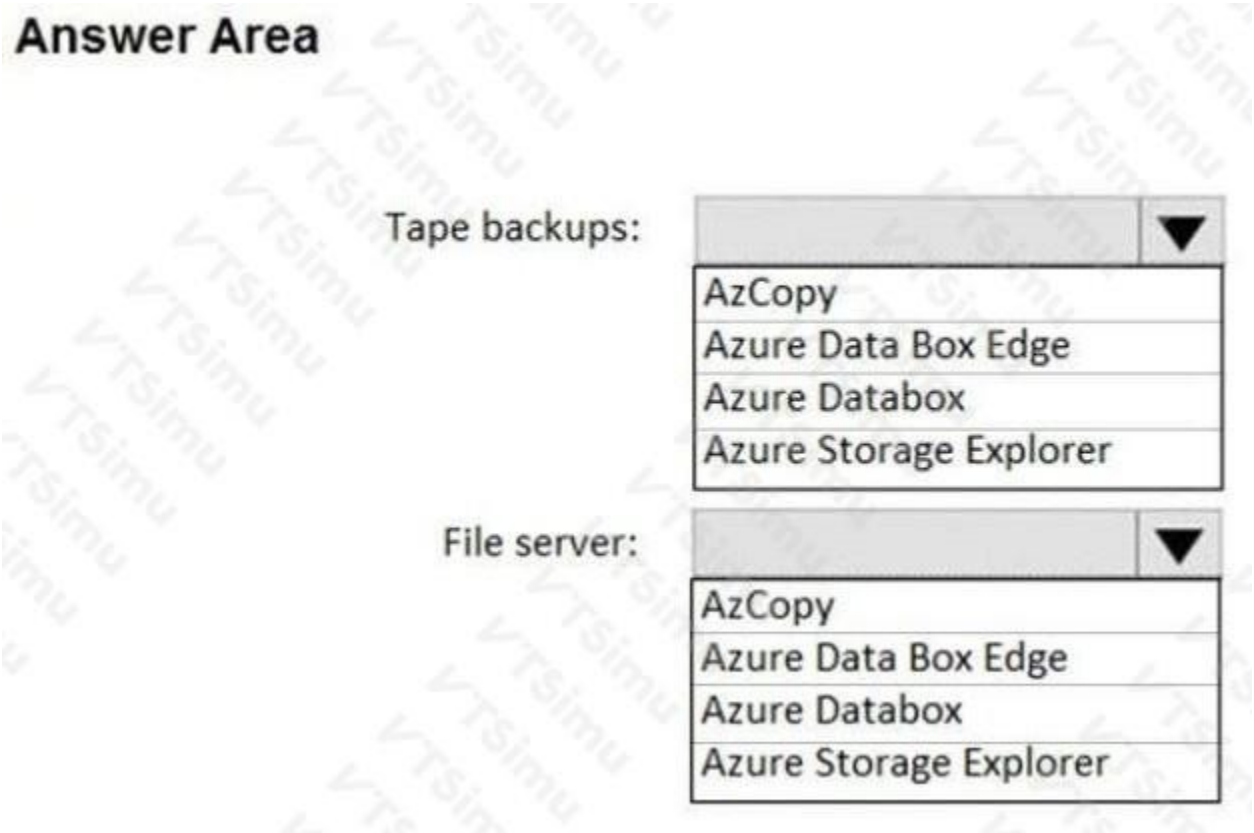

ANSWER:

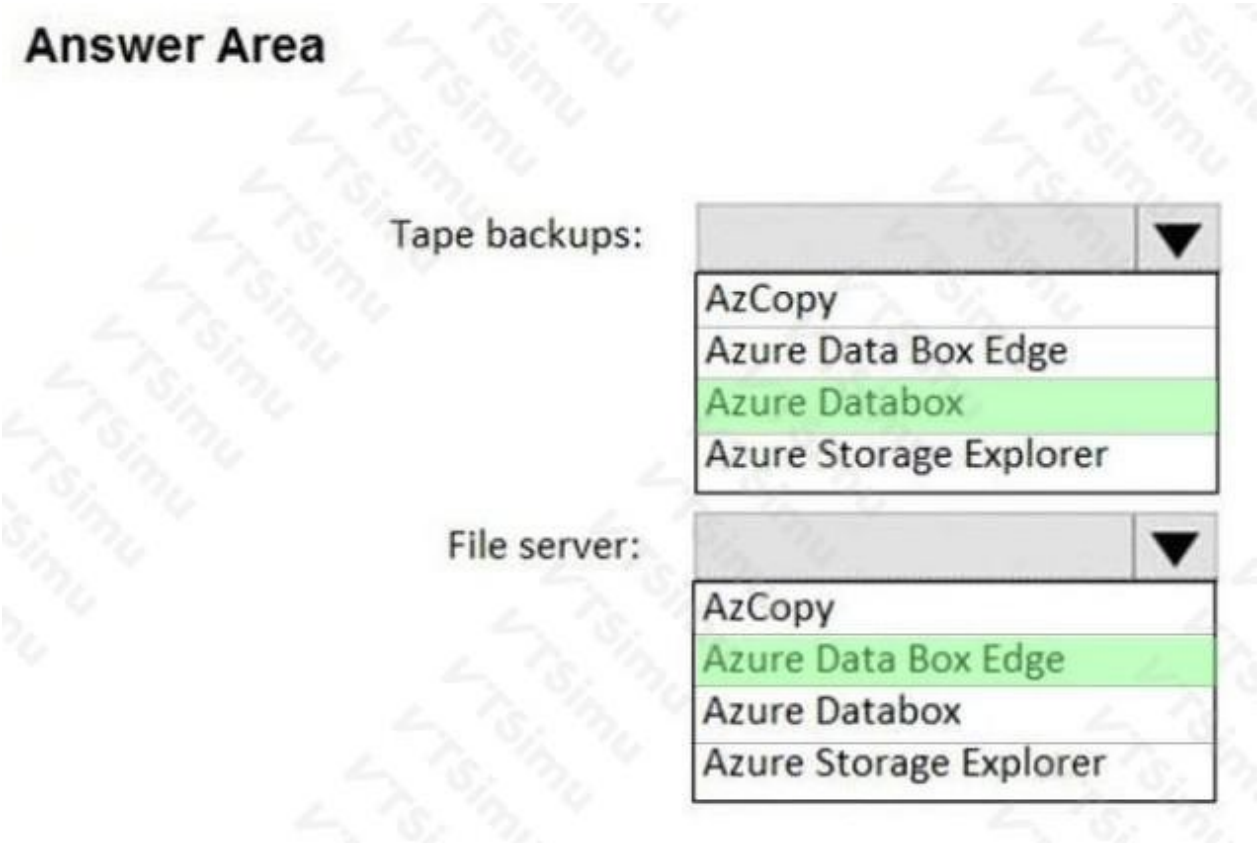

Explanation:

Use Azure Data Box for the tape backups and Azure Data Box Edge for the Windows file server. Azure Data Box is designed for large, offline bulk-data transfers when sending the data over the network would be slow or impractical. A Data Box device provides up to 80 TB of usable capacity, so it can accommodate the 50 TB of tape-backup data. After the backup data is copied to the encrypted device, the device is returned to Microsoft, where the data is uploaded into the selected Azure Storage destination. This makes it well suited to the one-time migration of a substantial tape-backup archive.

For the SAP Finance check images, Azure Data Box Edge is appropriate because the requirement is not merely an initial copy of 9 TB; it is continuous replication from the on-premises file server to Azure. Data Box Edge is a cloud-managed appliance that can operate as a storage gateway. Local systems write data to its file shares, and the appliance uploads that data to Azure Storage, enabling ongoing transfer as files are added or changed. This supports an initial migration followed by continuous synchronization of the image data without depending on a manual copy process. Microsoft's Data Box documentation describes Azure Data Box for offline bulk import and the Azure Stack Edge/Data Box Edge family for data transfer and gateway scenarios. See What is Azure Data Box? and What is Azure Stack Edge?.

QUESTION NO: 11

Your on-premises network contains a Microsoft SQL server named DCSQL1.

You have an Azure subscription that contains an SAP production landscape and a SQL server named AZSQL1. The SAP landscape uses DCSQLI and AZSQL1. The on-premises network is connected to Azure by using ExpressRoute.

You need to monitor network latency between A2SQL1 and DCSQL1 by using Connection Monitor. What should you install on AZSQL1?

A. the Azure Diagnostics extension

B. the Log Analytics agent

C. the Azure VM extension for SAP solutions

D. the Network Watcher extension

ANSWER: D

Explanation:

the Network Watcher extension is correct because Connection Monitor relies on the Azure Network Watcher Agent VM extension for monitoring connectivity and measuring network performance from an Azure virtual machine. The extension deploys the required agent components on AZSQL1, enabling it to act as a Connection Monitor endpoint. Connection Monitor can then collect measurements such as round-trip latency, packet loss, and connectivity status for traffic between the Azure SQL Server virtual machine and the on-premises DCSQL1 server over the ExpressRoute connection. This monitoring data is sent to the configured Log Analytics workspace, where it can be queried, visualized, and used to trigger alerts. Installing the Network Watcher extension on AZSQL1 is therefore the required Azure VM-side preparation for a Connection Monitor test involving that server. Microsoft documents the Network Watcher agent extension as the VM extension used by Connection Monitor, including deployment requirements for Azure VM endpoints. See Connection Monitor overview and Azure Network Watcher Agent virtual machine extension for Windows.

QUESTION NO: 12

You plan to deploy SAP NetWeaver ASCS and ERS instances on Linux virtual machines in separate Azure Availability Zones.

You need to provide a highly available virtual IP address for the ASCS service. The solution must support failover between the cluster nodes.

Which two Azure Load Balancer configurations should you include? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. An internal Azure Standard Load Balancer

B. A TCP health probe for the Pacemaker probe port

C. An Azure Basic Load Balancer

D. A public frontend IP address for the ASCS virtual IP

E. An Azure Application Gateway WAF policy

ANSWER: A B

Explanation:

You should use an **internal Azure Standard Load Balancer**. SAP ASCS is accessed internally by SAP application servers, and an internal frontend IP address provides the required private virtual IP address. The Standard SKU supports availability-zone architectures and is appropriate for production SAP high-availability deployments.

You must also configure a **health probe** on the probe port used by the Pacemaker azure-lb resource agent. The active cluster node responds to the probe, enabling the load balancer to direct traffic for the ASCS virtual IP address to that node after a failover.

For more information, see High availability for SAP NetWeaver on Azure and Azure Load Balancer overview.

QUESTION NO: 13

You plan to connect an on-premises SAP production landscape to Azure by using ExpressRoute.

You need to design the connection to minimize the effect of an ExpressRoute circuit failure and an Azure datacenter failure.

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

NOTE: Each correct selection is worth one point.

A. Provision ExpressRoute circuits in different peering locations.

B. Deploy a zone-redundant ExpressRoute virtual network gateway.

C. Deploy the ExpressRoute virtual network gateway in an Availability Set.

D. Use a single ExpressRoute circuit with a larger bandwidth SKU.

E. Configure an internal Azure Load Balancer for the gateway.

ANSWER: A B

Explanation:

Provisioning **ExpressRoute circuits in different peering locations** provides resiliency against a circuit, provider edge, or peering-location failure. The circuits can be connected to the Azure virtual network through ExpressRoute gateways and used with routing that supports failover between the connections.

A **zone-redundant ExpressRoute virtual network gateway** provides resiliency for the Azure gateway service across Availability Zones in supported regions. This protects the Azure-side gateway deployment from a zonal failure.

Using only one circuit does not protect against a circuit or peering-location failure. Similarly, an Availability Set does not provide zone resiliency for an ExpressRoute virtual network gateway.

For more information, see Designing for high availability with ExpressRoute and About ExpressRoute virtual network gateways.

QUESTION NO: 14 - (DRAG DROP)

A customer has an on-premises SAP environment.

The customer plans to migrate SAP to Azure.

You need to prepare the environment for the planned migration.

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

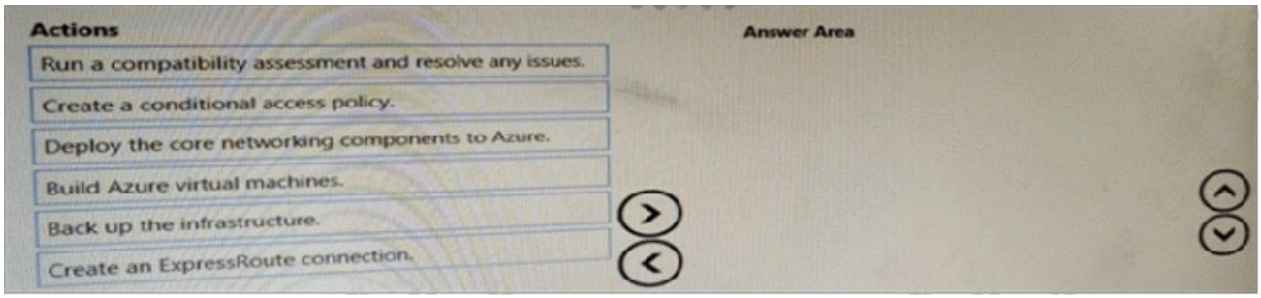

ANSWER:

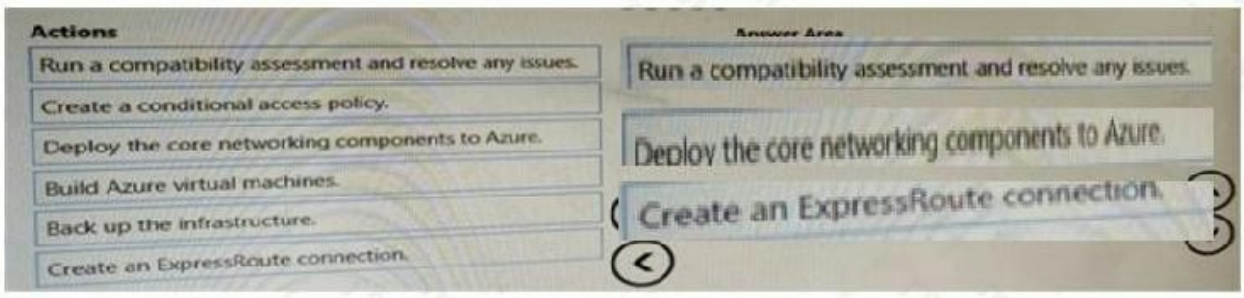

Explanation:

Start by running a compatibility assessment and resolving any issues. A migration plan for SAP needs a validated view of the existing landscape before resources are built in Azure. This assessment establishes whether the SAP applications, operating systems, database platform, interfaces, dependencies, and supporting components can be moved as planned. It also identifies remediation work, sizing requirements, and any constraints that must be addressed before migration activities begin. Microsoft's SAP on Azure guidance emphasizes assessing and planning the workload so that the target architecture and migration approach are based on the actual source environment.

Next, deploy the core networking components to Azure. The SAP target environment requires an Azure network foundation before it can support workload connectivity. This foundation includes the virtual network design, appropriate subnets, routing, security controls, name resolution considerations, and the connectivity components needed for hybrid access. Establishing these components first creates the landing-zone structure into which SAP systems and their private connectivity can be integrated.

Then create an ExpressRoute connection. ExpressRoute provides private, dedicated connectivity between the on-premises environment and Azure, which is well suited for enterprise SAP migration traffic and ongoing hybrid operations. Creating it after the Azure networking foundation allows the connection to be configured against the intended Azure network and gateway design. This sequence produces an assessed source environment, a ready Azure network, and then the private hybrid connection needed to support the planned migration. See Microsoft's production-ready SAP workload guidance and ExpressRoute overview for the relevant architecture and connectivity guidance.

QUESTION NO: 15 - (HOTSPOT)

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

NOTE: Each correct selection is worth one point.

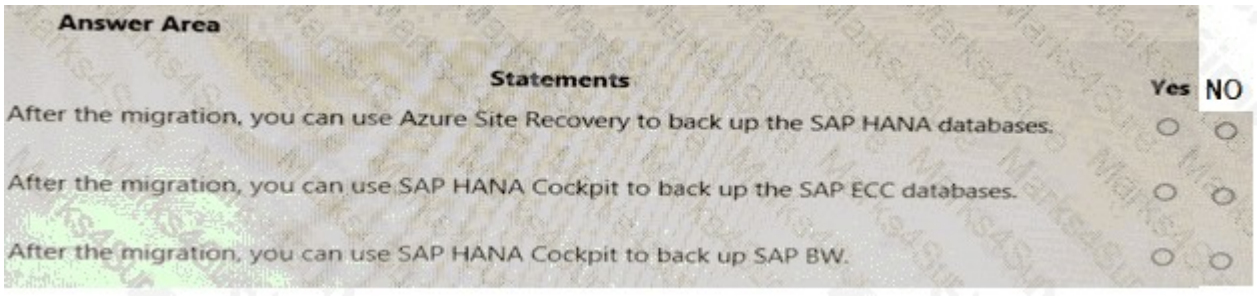

ANSWER:

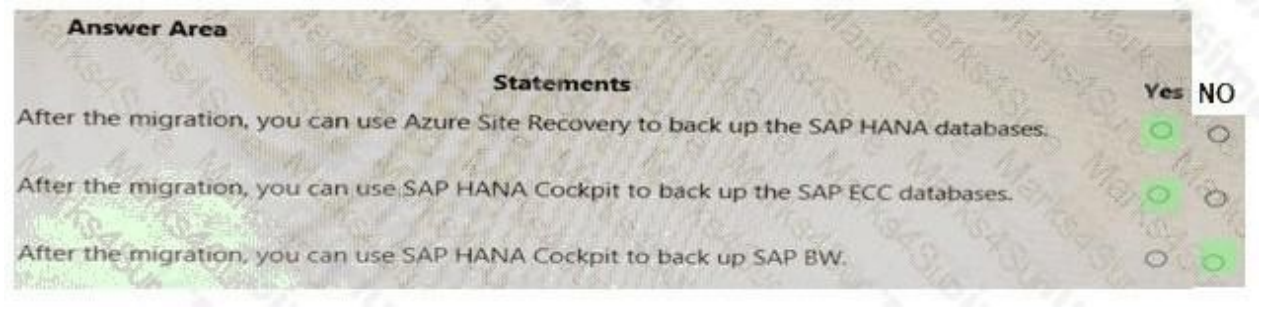

Explanation:

The correct selections are **No** for use of Azure Site Recovery to back up the SAP HANA databases, **Yes** for use of SAP HANA cockpit to back up the SAP ECC database, and **No** for use of SAP HANA cockpit to back up the SAP BW database.

Azure Site Recovery is designed to provide disaster recovery through replication, failover, and failback of supported workloads. It is not the SAP HANA database backup solution. A SAP HANA database backup must create usable database backup artifacts, including data backups and log backups, which are managed using SAP HANA's backup facilities and the appropriate backup destination. Microsoft distinguishes Azure Site Recovery disaster-recovery protection from backup protection in its Azure Site Recovery overview.

SAP HANA cockpit is the SAP administration interface for SAP HANA databases. It includes database-backup administration capabilities, such as initiating and monitoring backups and reviewing backup status, for the SAP HANA database that

supports the migrated SAP ECC system. SAP documents these functions in its SAP HANA cockpit backup and recovery documentation. This makes SAP HANA cockpit appropriate for administering the required HANA database backup activity for SAP ECC in the stated environment.

The selections therefore separate database backup administration from disaster-recovery replication and identify the HANA-backed SAP ECC workload as the one for which SAP HANA cockpit is used.

QUESTION NO: 16

You have an SAP production landscape in Azure.

Users access the landscape from the internet by using Azure Application Gateway.

You need to analyze the network traffic of the landscape by using Azure Network Watcher traffic analytics.

What should you configure?

A. network security group (NSG) flow logs

B. Network Manager

C. Connection Monitor

D. Diagnostic settings in Application Gateway

ANSWER: A

Explanation:

network security group (NSG) flow logs is correct because Traffic Analytics relies on flow-log data to provide aggregated visibility into network communications. NSG flow logs record IP traffic that is allowed or denied by network security groups, including source and destination addresses, ports, protocol, flow direction, and traffic volume. After the logs are sent to an Azure Storage account, Traffic Analytics processes them and presents insights in Azure Monitor, such as traffic patterns, application-protocol information, potentially malicious flows, and traffic-rule usage. This enables administrators to investigate network activity associated with the SAP landscape and identify communication trends or anomalous traffic while maintaining the network controls that protect workloads behind Application Gateway. The relevant NSGs must be associated with the appropriate network interfaces or subnets, and their flow logs must be enabled for the workload traffic that requires analysis. Microsoft documents Traffic Analytics as a solution that analyzes network security group flow logs. See Traffic Analytics documentation and NSG flow logs overview.

QUESTION NO: 17

You have an on-premises SAP NetWeaver development landscape that contains the resources shown in the following table.

Name	Description
SAPDB1	Hyper-V virtual machine that runs Microsoft SQL Server 2017 and contains a 30-TB database
SAPSRV1	Hyper-V virtual machine that runs Windows Server

You have a 500-Mbps ExpressRoute circuit between the on-premises environment and a virtual network.

You plan to migrate the landscape to Azure.

What should you include in the solution?

A. Azure Site Recovery

B. Microsoft System Center 2019 - Data Protection Manager (DPM 2019)

C. Azure Data Box

D. Azure Backup Server

ANSWER: A

Explanation:

Azure Site Recovery is appropriate because it can replicate supported on-premises virtual machines to Azure and orchestrate a controlled migration through test failovers and a final planned failover. For a SAP NetWeaver development landscape, this provides a practical migration path that uses the existing ExpressRoute connectivity for replication traffic while allowing the environment to be validated in Azure before the production migration cutover. Site Recovery continuously replicates machine data to Azure-managed disks, so the target virtual machines can be created and started in Azure with the replicated operating system, application, and database state. A planned failover helps minimize data loss by synchronizing remaining changes before the on-premises machines are shut down and Azure machines are brought online. This approach is particularly useful when the goal is to migrate the existing landscape rather than redesigning or restoring it from backups. Microsoft documents Azure Site Recovery as a migration solution for replicating VMware virtual machines and physical servers to Azure, including planned failover capabilities. For SAP-specific guidance, Microsoft also describes using Site Recovery to migrate SAP deployments to Azure. See Migrate machines to Azure with Site Recovery and Migrate an SAP workload to Azure.

QUESTION NO: 18

You are planning a deployment of SAP on Azure that will use SAP HANA.

You need to ensure that the SAP application servers are in the same datacenter as the HANA nodes.

What should you use?

A. a resource group

B. a virtual machine scale set
an application group

C. a proximity placement group

ANSWER: C

Explanation:

a proximity placement group is the appropriate choice because it is an Azure logical grouping used to keep compute resources physically close together within a single Azure datacenter. For an SAP deployment, placing the SAP application-server virtual machines and SAP HANA virtual machines in the same proximity placement group helps minimize network latency between the application tier and the database tier. Low and predictable latency is particularly important for SAP workloads because application servers issue frequent database requests to SAP HANA, and response time can affect overall transactional performance.

A proximity placement group expresses a colocation intent to Azure when virtual machines are deployed. Azure attempts to place the members as near as possible, typically in the same datacenter, rather than allowing them to be distributed across different datacenters in a region. The proximity placement group should be created before the associated virtual machines and specified during their deployment. Capacity must also be available in the intended datacenter for the required VM sizes. Microsoft documents proximity placement groups and their use for latency-sensitive workloads at Azure virtual machine proximity placement groups. SAP-specific infrastructure guidance is available in SAP on Azure documentation.

QUESTION NO: 19

Your company has a an on-premises SAP environment.

Recently, the company split into two companies named Litware, inc and Contoso.Ltd. Litware retained the SAP environment.

Litware plans to export data that is relevant only to Contoso. The export will be 1.5 TB.

Contoso build a new SAP environment on Azure.

You need to recommend a solution for Litware to make the data available to Contoso in Azure. The solution must meet the following requirements:

Minimize the impact on the network.

Minimize the administrative effort for Litware.

What should you include in the recommendation?

A. Azure Migrate

B. Azure Databox

C. Azure Site Recovery

D. Azure import/Export service

ANSWER: B

Explanation:

Azure Databox is the appropriate recommendation because it transfers data to Azure by using a Microsoft-provided physical device rather than sending the full 1.5 TB through Litware's network connection. Litware copies the exported SAP data to the device locally and ships it to Microsoft. After Microsoft receives the device, the data is uploaded into the designated Azure Storage account, where Contoso can access it for its new Azure-based SAP environment. This approach substantially minimizes network impact, since the bulk data movement is performed offline. It also minimizes Litware's administrative effort compared to an import/export workflow that requires the customer to procure, prepare, encrypt, and ship its own compatible storage drives. Azure Data Box is designed for secure, offline transfer of large volumes of data when network transfer is impractical, slow, or undesirable. The service includes device handling, secure data erasure after ingestion, and portal-based tracking of the transfer job. Although the export size is modest relative to the appliance capacity, Azure Data Box remains the best fit among the listed choices because it directly satisfies both the low-network-impact and low-administration requirements. See Azure Data Box overview and Deploy Azure Data Box.

QUESTION NO: 20

You are deploying SAP on Azure. The database server will use SAP HAN

A.
The application servers will run Windows Server. You need to test network latency and throughput between the frontend SAP servers and the database servers. Which three tools can you use to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. SockPerf

B. HCMT

C. NIPING

D. Latte.exe

E. IOMeter

ANSWER: A C D

Explanation:

SockPerf, **NIPING**, and **Latte.exe** are appropriate utilities for assessing network behavior relevant to SAP application-to-database communication. SockPerf is a network benchmarking tool that measures latency and can generate high-rate traffic to assess throughput characteristics. It is particularly useful from Linux-based environments, including SAP HANA database hosts. NIPING is SAP's network interface test utility and is designed to validate SAP communication paths by measuring connection behavior, latency, and transfer performance between systems running SAP components. It is available for the

operating systems commonly used by SAP application and database tiers. Latte.exe is a Microsoft network testing utility for Windows that can measure TCP latency and throughput, making it useful when testing from Windows Server-based SAP frontend or application servers. Together, these tools provide practical measurements of the network path rather than relying solely on nominal virtual-network specifications. Microsoft recommends measuring actual network performance between virtual machines because throughput and latency can vary by VM size, placement, operating system configuration, and traffic pattern. See Azure virtual network bandwidth testing and Test virtual network latency.

QUESTION NO: 21 - (SIMULATION)

You have an SAP production landscape that uses SAP HANA databases.

You configure a metric alert for the primary HANA server as shown in the following exhibit.

Answer Area

Statements	Yes	No
HANA Admins will be alerted by email if the server is at 85 percent for one minute, and then lowers to 40 percent.	○	○
HANA Admins will be alerted if the server is at 95 percent for 15 minutes.	○	○
Amy@contoso.com will be alerted by email if the server CPU cycles between 80 and 90 percent for 15 minutes.	○	○

ANSWER: See the explanation for the answer

Explanation:

Answers:

The alert evaluates the CPU signal over a **15-minute** period. Therefore, a one-minute CPU value of 85 percent does not satisfy the configured evaluation duration. A CPU value of 95 percent for the whole 15-minute period meets the threshold. CPU values between 80 and 90 percent also meet the configured CPU threshold for the required period, and the configured email action notifies Amy.

QUESTION NO: 22

You are deploying an SAP environment on Azure that will use an SAP HANA database server.

You provision an Azure virtual machine for SAP HANA by using the M64s virtual machine SKU.

You need to set the swap space by using the Microsoft Azure Linux Agent (waagent) configuration file.

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

NOTE: Each correct selection is worth one point.

A. ResourceDisk.EnableSwapEncryption=n

B. AutoUpdate.Enabled=n

C. ResourceDisk.SwapSizeMB=229376

D. ResourceDisk.EnableSwap=y

ANSWER: C D

Explanation:

`ResourceDisk.EnableSwap=y` enables swap-file creation by the Azure Linux Agent during provisioning. The agent creates the file on the VM resource disk at the configured resource-disk mount point, allowing swap configuration to be applied consistently when the VM is provisioned or the agent processes its configuration.

`ResourceDisk.SwapSizeMB=229376` specifies the required swap-file capacity in megabytes for this SAP HANA deployment on an M64s VM. The value is 224 GiB and supplies the sizing specified for this workload. The size setting alone does not create swap; it works together with swap enablement. Therefore, both settings are required in `/etc/waagent.conf` to define and activate the intended swap space.

The Azure Linux Agent configuration supports the `ResourceDisk.EnableSwap` and `ResourceDisk.SwapSizeMB` parameters for managing a swap file on the resource disk. After changing agent configuration, apply the change according to the distribution's Azure Linux Agent service-management procedure. See Azure Linux Agent overview and SAP HANA on Azure VM operations.

QUESTION NO: 23

You deploy an SAP landscape on Azure.

You plan to use an Azure Automation account to stop the SAP virtual machines outside of business hours.

You need to ensure that you can use Azure Automation runbook a for the virtual machines.

What does the Azure Automation account require?

A. an Azure Storage account

B. an Azure App Service WebJob

C. a Recovery Services vault

D. a managed identity

ANSWER: D

Explanation:

a managed identity is required so that the Azure Automation runbook can authenticate to Azure Resource Manager and perform management operations on the SAP virtual machines, such as stopping and starting them. A system-assigned managed identity is created for the Automation account and is automatically managed by Azure, avoiding the need to store credentials, certificates, or service-principal secrets in the runbook. The identity must then receive Azure RBAC permissions at the appropriate scope, such as the resource group containing the SAP virtual machines. For example, the Virtual Machine Contributor role allows the identity to perform VM lifecycle actions while following least-privilege principles. In a PowerShell runbook, the managed identity can be used with commands such as `Connect-AzAccount -Identity`, after which Azure VM cmdlets can be used to stop the required machines. This is the current supported approach for authenticating Azure Automation runbooks. Azure Automation Run As accounts were retired and should not be used for new implementations. See Use managed identities with Azure Automation and Run As account retirement.

QUESTION NO: 24 - (DRAG DROP)

DRAG DROP

You migrate SAP ERP Central Component (SAP ECC) production and non-production landscapes to Azure.

You are licensed for SAP Landscape Management (LaMa).

You need to refresh from the production landscape to the non-production landscape.

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

Select and Place:

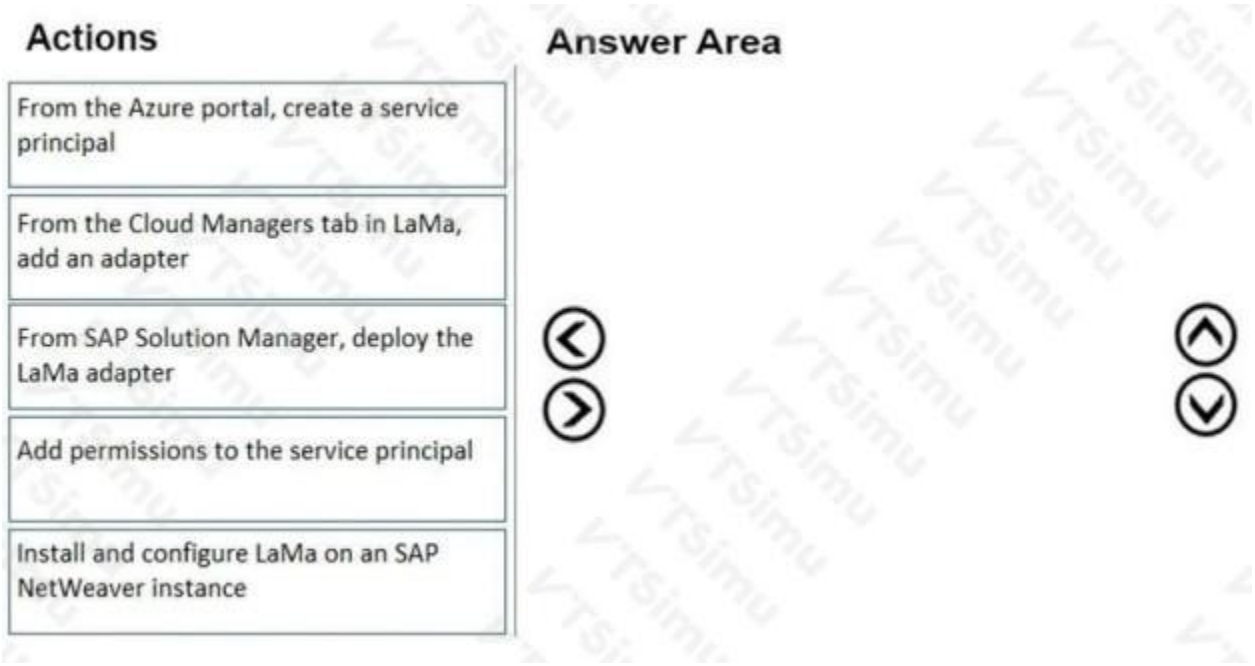

ANSWER:

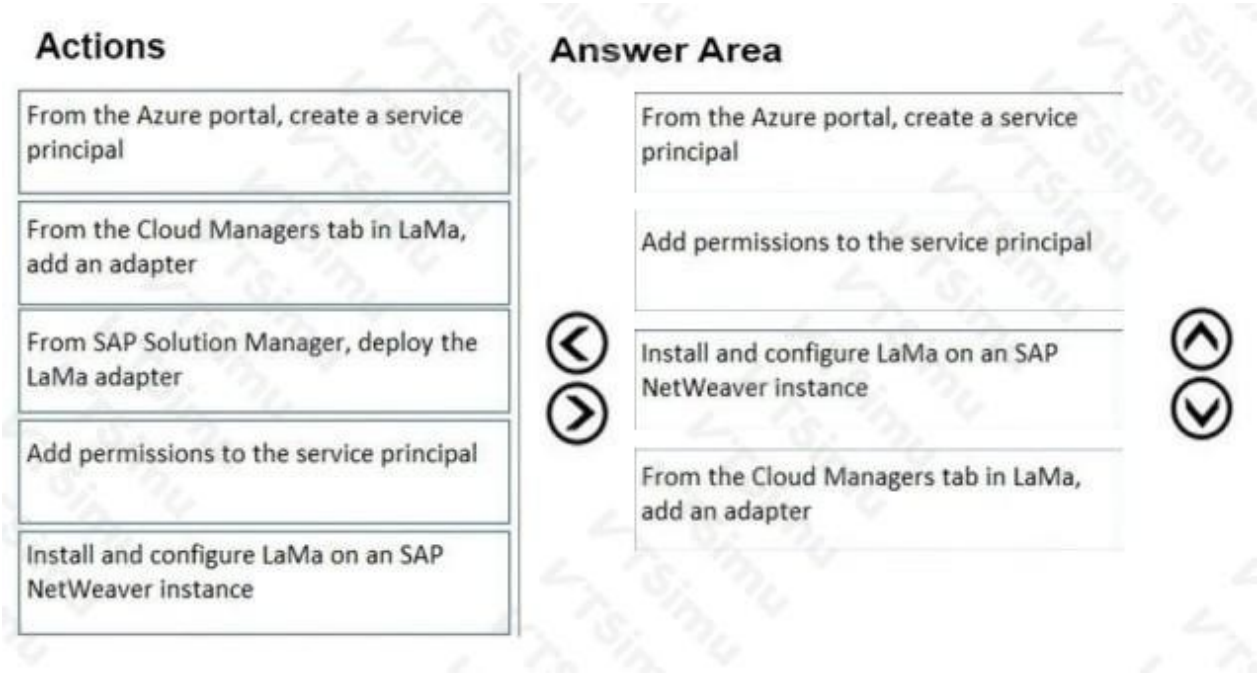

Explanation:

To prepare SAP Landscape Management to perform a production-to-non-production refresh in Azure, first create a service principal in the Azure portal. The Azure connector used by SAP Landscape Management needs a Microsoft Entra application identity for authenticated access to Azure APIs. After creating that identity, assign it the necessary Azure permissions at the required scope so it can manage the relevant Azure resources during SAP system operations, such as provisioning or controlling virtual machines and related infrastructure.

Next, install and configure SAP Landscape Management on an SAP NetWeaver instance. This establishes the SAP Landscape Management application and its administration interface. Once the SAP Landscape Management system is available and configured, use its Cloud Managers tab to add the Microsoft Azure cloud adapter. Adding the adapter at this

point connects the already installed SAP Landscape Management environment to Azure by using the previously created and authorized service principal. The adapter provides the integration SAP Landscape Management needs to orchestrate Azure infrastructure as part of lifecycle tasks, including system copies and refresh-related operations.

This order ensures that the Azure authentication identity exists before it is referenced by the connector, and that the SAP Landscape Management UI and infrastructure are in place before an administrator configures a Cloud Manager. Microsoft's SAP guidance describes the Azure integration and SAP Landscape Management deployment prerequisites in the SAP Landscape Management installation guidance. For Azure role assignment concepts used to grant the service principal access, see Assign Azure roles using the Azure portal.

QUESTION NO: 25

You migrate an SAP environment to Azure.

You need to inspect all the outbound traffic from the SAP application servers to the Internet.

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

NOTE: Each correct selection is worth one point.

A. Azure Traffic Manager

B. Azure Firewall

C. Network Performance Monitor

D. Azure user-defined routes

E. Azure Load Balancer NAT rules

F. a Web Application Firewall (WAF) for Azure Application Gateway

ANSWER: B D

Explanation:

Azure Firewall and Azure user-defined routes provide the required centralized egress-inspection design. Deploy Azure Firewall in a dedicated subnet and configure its network and application rules to evaluate outbound connections from the SAP application-server subnets. Azure Firewall is a stateful, fully managed firewall that can control and log outbound network traffic, including traffic filtered by destination, port, protocol, and—when application rules are applicable—fully qualified domain name.

Azure user-defined routes are required to ensure that the application servers' Internet-bound traffic actually traverses the firewall. A route table associated with each SAP application-server subnet can define a default route for `0.0.0.0/0` whose next hop is the firewall's private IP address, using the virtual-appliance next-hop type. This forces outbound flows through Azure Firewall rather than allowing direct Internet egress. Firewall diagnostic logs can then be sent to Log Analytics, a storage account, or an event hub for monitoring, audit, and investigation. This routing-and-firewall combination is the standard Azure pattern for inspecting and governing outbound traffic from workloads in virtual networks. See Deploy and configure Azure Firewall and User-defined routes in Azure.

QUESTION NO: 26

You plan to deploy an SAP environment on Azure. The SAP environment will have landscapes for production, development. and quality assurance.

You need to minimize the costs associated with running the development and quality assurance landscapes on Azure. What should you do?

A. Create Azure Automation runbooks to stop, deallocate, and start Azure virtual machines.

B. Create a scheduled task that runs the stopsap command.

C. Configure scaling for Azure App Service.

D. Configure Azure virtual machine scales sets.

ANSWER: A

Explanation:

Create Azure Automation runbooks to stop, deallocate, and start Azure virtual machines. Development and quality-assurance SAP systems commonly have predictable periods when they are not needed, such as evenings and weekends. An Azure Automation runbook can schedule a controlled shutdown of the SAP workload and then deallocate its virtual machines. Deallocation is essential for cost reduction because it releases the compute allocation; charges for the virtual-machine compute resources stop while the machines are deallocated. The runbook can also start the environment before the next working period, allowing teams to use the systems without a manual startup process.

For SAP landscapes, the automation should implement an application-aware sequence, stopping SAP services and dependent components cleanly before deallocating the VMs, and starting infrastructure and SAP components in their required order. Azure Automation provides runbooks for automating these repeatable operational tasks and supports scheduling their execution. This approach retains the development and quality-assurance environments when needed while avoiding unnecessary compute spend during idle periods. See Azure Automation runbooks documentation and Microsoft's guidance on automating SAP system start and stop.

QUESTION NO: 27

You plan to deploy an SAP environment on Azure. The SAP environment will have landscapes for production, development and quality assurance.

You need to minimize the costs associated with running the development and quality assurance landscapes on Azure.

What should you do?

A. Configure scaling for Azure App Service

B. Create a scheduled task that runs the stopsap command

C. Configure Azure virtual machine scale sets

D. Create Azure Automation runbooks to stop, deallocate, and start Azure virtual machines

ANSWER: D

Explanation:

Creating Azure Automation runbooks to stop, deallocate, and start Azure virtual machines is the appropriate way to reduce costs for nonproduction SAP development and quality-assurance landscapes. These systems commonly do not need to be available continuously, so runbooks can schedule an orderly SAP shutdown outside business hours, stop the associated Azure virtual machines, and deallocate them. Deallocation is essential because it releases the compute allocation; compute charges for the virtual machines cease while they are deallocated, although charges for retained resources such as managed disks and storage continue. The same automation can start the virtual machines before working hours and coordinate the required system startup sequence, such as database, central services, and application servers. Azure Automation supports scheduled runbook execution and can use identities and Azure management tooling to perform these actions consistently without manual operational effort. This approach balances cost optimization with the operational requirement to make development and QA systems available when teams need them. Microsoft specifically documents using Azure Automation to automate SAP system start and stop operations, and Azure documents the billing distinction between stopped and deallocated virtual machines. See Automate SAP system start and stop with Azure Automation and Azure virtual machine states and billing status.

QUESTION NO: 28

You deploy SAP HANA databases on Azure virtual machines.

You need to use Azure Backup to protect the SAP HANA databases. The solution must provide application-consistent database backups that are managed from a Recovery Services vault.

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

NOTE: Each correct selection is worth one point.

A. a Recovery Services vault

B. the Backint agent for Azure Backup

C. an Azure Site Recovery vault

D. an Azure Backup Server installation

E. an Azure File Sync server endpoint

ANSWER: A B

Explanation:

A **Recovery Services vault** is required to store and manage the backup configuration, recovery points, retention policy, and restore operations for SAP HANA databases protected by Azure Backup.

The **Backint agent for Azure Backup** is also required on the SAP HANA database host. Azure Backup uses the SAP HANA Backint interface to perform database-aware backup and restore operations. The agent transfers the backup data created by SAP HANA to the Recovery Services vault.

For more information, see About SAP HANA database backup in Azure Backup and Back up SAP HANA databases in Azure.

QUESTION NO: 29

You migrate an on-premises instance of SAP NANA that runs SUSE Linux Enterpnse Server (SLES) to an Azure virtual machine.

You project that in two years, you will replace the virtual machine with a larger virtual machine within the same flexibility group.

You need to recommend solutions to minimize HANA deployment costs during the next three years. The solutions must not affect the availability SLAs.

Which two solutions should you recommend? Each correct answer presents a complete solution

NOTE: Each correct selection is worth one point.

A. a three-year reservation that has instance size flexibility

B. a one-year reservation that has instance size flexibility

C. a one-year reservation that has capacity priority

D. Azure Hybrid Benefit

E. Azure Spot instance

ANSWER: A D

Explanation:

A three-year reservation that has instance size flexibility is appropriate because the workload is expected to run for the full three-year period and will move to a larger VM after two years. Azure VM reservations provide a significant discount compared to pay-as-you-go pricing. Instance size flexibility allows the reservation benefit to apply across eligible VM sizes in

the same instance-size flexibility group, so changing to the planned larger VM does not require abandoning the reservation or losing its pricing benefit. Reservations are a billing discount and do not change the virtual machine's availability architecture or service-level agreement.

Azure Hybrid Benefit is also appropriate for SUSE Linux Enterprise Server when the organization has eligible SUSE Linux subscriptions. It enables existing qualifying subscription value to be used for SLES virtual machines in Azure, reducing the operating-system software cost while retaining the Azure VM configuration and associated availability characteristics. The benefit can be combined with the VM reservation because they address separate portions of cost: the reservation reduces Azure compute charges, while Azure Hybrid Benefit reduces eligible Linux software charges. Microsoft describes reservation instance-size flexibility at Azure VM reservation instance size flexibility and eligibility for SLES at Azure Hybrid Benefit for Linux.

QUESTION NO: 30

You plan to deploy an SAP environment on Azure.

You plan to store all SAP connection strings securely in Azure Key Vault without storing credentials on the

Azure virtual machines that host SAP.

What should you configure to allow the virtual machines to access the key vault?

A. Azure Active Directory (Azure AD) Privilege Identity Manager (PIM)

B. role-based access control (RBAC)

C. a Managed Service Identity (MSI)

D. the Custom Script Extension

ANSWER: C

Explanation:

A Managed Service Identity (MSI), now generally called a managed identity, is the appropriate configuration because it gives an Azure virtual machine an identity in Microsoft Entra ID without requiring passwords, certificates, client secrets, or other credentials to be stored on that VM. The SAP application or supporting scripts running on the virtual machine can request an access token from Azure by using that identity, then use the token to authenticate to Azure Key Vault and retrieve the required connection-string secrets.

For this design, enable either a system-assigned managed identity on each SAP virtual machine or assign a user-assigned managed identity to the relevant virtual machines. Then grant that identity only the required Key Vault data-plane permissions, preferably through an Azure RBAC role such as Key Vault Secrets User when the vault uses the Azure RBAC permission model. This creates a credential-free authentication path for the workload while keeping the secret values centralized in Key Vault. Microsoft documents managed identities as the mechanism for Azure resources to authenticate to services that support Microsoft Entra authentication, including Key Vault. See Managed identities for Azure resources and Authenticate to Azure Key Vault.

QUESTION NO: 31

You have an Azure subscription that contains an SAP HANA on Azure (Large Instances) deployment. The deployment is forecasted to require an additional 256 GB of storage. What is the minimum amount of additional storage you can allocate?

A. 256 GB

B. 512 GB

C. 1TB

D. 2 TB

ANSWER: C

Explanation:

1TB is correct because SAP HANA on Azure (Large Instances) storage expansion is provisioned in fixed minimum increments of 1 TB. Although the deployment's forecast identifies a 256-GB capacity shortfall, Azure cannot allocate only that fractional amount for this platform's storage expansion. The requested increase must therefore be rounded up to the smallest supported allocation unit, which is 1 TB. This behavior reflects the dedicated, high-performance storage architecture used by HANA Large Instances, where capacity changes are planned and provisioned as larger storage increments rather than as arbitrary per-GB additions. When planning the increase, account for not only the immediate additional data requirement but also SAP HANA operational headroom, including space needed for data growth, log volumes, backups, snapshots, and maintenance activities. Capacity planning should also consider the throughput and IOPS implications of the storage configuration, rather than treating storage solely as a raw capacity requirement. Microsoft's SAP HANA on Azure (Large Instances) documentation describes the platform architecture and its storage characteristics, while the operations guidance covers managing and planning the environment. See SAP HANA on Azure (Large Instances) architecture and SAP HANA on Azure (Large Instances) operations.

QUESTION NO: 32

You plan to migrate an SAP ERP Central Component (SAP ECC) production system to Azure.

You are reviewing the SAP EarlyWatch Alert report for the system.

You need to recommend sizes for the Azure virtual machines that will host the system.

Which two sections of the report should you review? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Hardware Capacity

B. Patch Levels under SAP Software Configuration

C. Hardware Configuration under Landscape

D. Database and ABAP Load Optimization

E. Data Volume Management

ANSWER: A C

Explanation:

Hardware Capacity is the primary SAP EarlyWatch Alert section for evaluating whether the current system has sufficient compute resources and for identifying workload-driven sizing requirements. Its capacity checks assess resource consumption trends and peak demand, including processor, memory, database, and throughput-related indicators. For an Azure migration, these measurements provide the evidence needed to translate the production workload into required SAPS, virtual CPU capacity, and memory for the target virtual machines, with appropriate headroom for future growth.

Hardware Configuration under Landscape supplies the complementary baseline of the current SAP landscape. It documents the existing hosts and their hardware characteristics, such as processor and memory configuration, which helps correctly interpret the capacity findings and map individual SAP application and database roles to Azure virtual machine sizes. Reviewing both sections prevents sizing based solely on a static server inventory or solely on utilization data; the recommendation instead reflects the deployed landscape and its measured production demand. SAP describes the capacity-oriented EarlyWatch Alert checks in its Hardware Capacity Checks in EWA documentation. For Azure-specific sizing principles, see Azure SAP workload sizing and benchmarking guidance.

QUESTION NO: 33 - (SIMULATION)

You deploy an SAP environment on Azure.

You need to grant an SAP administrator read-only access to the Azure subscription. The SAP administrator must be prevented from viewing network information.

How should you configure the role-based access control (RBAC) role definition? To answer, drag the appropriate values to the correct targets. 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.

NOTE: Each correct selection is worth one point.

ANSWER: See the explanation for the answer

Explanation:

Create a custom RBAC role at the subscription scope with the following permissions:

```
{
  "Actions": [
    "*/read"
  ],
  "NotActions": [
    "Microsoft.Network/*/read"
  ],
  "AssignableScopes": [
    "/subscriptions/<subscription-ID>"
  ]
}
```

`*/read` grants read-only access to Azure resource types.
`Microsoft.Network/*/read` in `NotActions` excludes all Microsoft.Network read operations, preventing the administrator from viewing network resources/information.
Set `AssignableScopes` to the target Azure subscription.

QUESTION NO: 34

You plan to use Azure Monitor for SAP solutions to monitor an SAP landscape that contains SAP HANA, SAP NetWeaver, and Linux virtual machines.

You need to collect SAP-specific telemetry and create operational alerts from the collected data.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure Monitor for SAP solutions resource.

B. Configure SAP HANA and SAP NetWeaver provider configurations.

C. Deploy an Azure Traffic Manager profile.

D. Create an Azure DNS private resolver.

E. Enable Azure DDoS Protection Standard.

ANSWER: A B

Explanation:

Create an **Azure Monitor for SAP solutions resource** for the SAP landscape. This resource is the monitoring service used to configure SAP-specific providers and visualize SAP workload telemetry through Azure Monitor.

You must also configure the required **provider configurations**, such as SAP HANA and SAP NetWeaver providers, and associate the monitoring resource with a Log Analytics workspace. The provider configurations collect the applicable SAP metrics, which can then be queried and used by Azure Monitor alert rules.

For more information, see What is Azure Monitor for SAP solutions? and Configure alerts for Azure Monitor for SAP solutions.

QUESTION NO: 35

You manage SAP production landscapes in several Azure subscriptions.

You need to prevent deployment of virtual machines that are not approved for SAP workloads. The restriction must be applied consistently to all the subscriptions.

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

NOTE: Each correct selection is worth one point.

A. an Azure management group

B. an Azure Policy assignment that restricts allowed virtual machine SKUs

C. an Azure resource lock on each subscription

D. an Azure RBAC assignment for Virtual Machine Contributor

E. an Azure Availability Set

ANSWER: A B

Explanation:

An **Azure management group** provides a governance scope above subscriptions. Assignments made at this scope are inherited by the subscriptions and their contained resource groups, which enables a consistent control across the SAP production subscriptions.

An **Azure Policy assignment** that uses the Allowed virtual machine SKUs policy definition can restrict deployment to a defined list of approved Azure virtual machine sizes. This helps ensure that administrators deploy only the VM SKUs selected and validated for the SAP workload.

Azure RBAC controls who can create or manage resources, but it does not evaluate whether a requested virtual machine SKU is approved. Resource locks prevent changes or deletion of existing resources and do not enforce deployment standards.

For more information, see Azure management groups overview and Azure Policy built-in definitions.