

➤ **Vendor: Microsoft**

➤ **Exam Code: AZ-120**

➤ **Exam Name: Planning and Administering Microsoft Azure for SAPWorkloads**

➤ **New Updated Questions from [Braindump2go](https://www.braindump2go.com) (Updated in [August/2020](#))**

[Visit Braindump2go and Download Full Version AZ-120 Exam Dumps](#)

QUESTION 22

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

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

You deploy SAP HANA on Azure (Large Instances).

You need to back up the SAP HANA database to Azure.

Solution: You create a Recovery Services vault and a backup policy.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Backup architecture

The backup process begins by creating a Recovery services vault in Azure. This vault will be used to store the backups and recovery points created over time.

The Azure VM running SAP HANA server is registered with the vault, and the databases to be backed-up are discovered. To enable the Azure Backup service to discover databases, a preregistration script must be run on the HANA server as a root user.

This script creates AZUREWLBACKUPHANAUSER DB user and a corresponding key with the same name in hdbuserstore. Refer to the setting up permissions section to understand more about what the script does.

Azure Backup Service now installs the Azure Backup Plugin for HANA on the registered SAP HANA server.

The AZUREWLBACKUPHANAUSER DB user created by the preregistration script is used by the Azure Backup Plugin for HANA to perform all backup and restore operations. If you attempt to configure backup for SAP HANA DBs without running this script, you might receive the following error:

UserErrorHanaScriptNotRun.

To configure backup on the databases that are discovered, choose the required backup policy and enable backups.

Once the backup is configured, Azure Backup service sets up the Backint parameters at the DATABASE level on the protected SAP HANA server.

The Azure Backup Plugin for HANA maintains all the backup schedules and policy details. It triggers the scheduled backups and communicates with the HANA Backup Engine through the Backint APIs.

The HANA Backup Engine returns a Backint stream with the data to be backed up.

All the scheduled backups and on-demand backups (triggered from the Azure portal) that are either full or differential are initiated by the Azure Backup Plugin for HANA. However, log backups are managed and triggered by HANA Backup Engine itself.

References:

[AZ-120 Exam Dumps](#) [AZ-120 Exam Questions](#) [AZ-120 PDF Dumps](#) [AZ-120 VCE Dumps](#)

<https://www.braindump2go.com/az-120.html>

<https://docs.microsoft.com/en-us/azure/backup/sap-hana-db-about>

<https://docs.microsoft.com/en-us/azure/backup/backup-azure-sap-hana-database#configure-backup>

QUESTION 23

You have an SAP environment on Azure that uses multiple subscriptions.

To meet GDPR requirements, you need to ensure that virtual machines are deployed only to the West Europe and North Europe Azure regions.

Which Azure components should you use?

- A. Azure resource locks and the Compliance admin center
- B. Azure resource groups and role-based access control (RBAC)
- C. Azure management groups and Azure Policy
- D. Azure Security Center and Azure Active Directory (Azure AD) groups

Answer: C

Explanation:

Azure Policy enables you to set policies to conform to the GDPR. Azure Policy is generally available today at no additional cost to Azure customers. You can use Azure Policy to define and enforce policies that help your cloud environment become compliant with internal policies as well as external regulations.

Azure Policy is deeply integrated into Azure Resource Manager and applies across all resources in Azure.

Individual policies can be grouped into initiatives to quickly implement multiple rules. You can also use Azure Policy in a wide range of compliance scenarios, such as ensuring that your data is encrypted or remains in a specific region as part of GDPR compliance. Microsoft is the only hyperscale cloud provider to offer this level of policy integration built in to the platform for no additional charge.

References:

<https://azure.microsoft.com/de-de/blog/new-capabilities-to-enable-robust-gdpr-compliance/>

QUESTION 24

A customer that has a large enterprise SAP environment plans to migrate to Azure. The environment uses servers that run Windows Server 2016 and Microsoft SQL Server.

The environment is critical and requires a comprehensive business continuity and disaster recovery (BCDR) strategy that minimizes the recovery point objective (RPO) and the recovery time objective (RTO).

The customer wants a resilient environment that has a secondary site that is at least 250 kilometers away.

You need to recommend a solution for the customer.

Which two solutions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. warm standby virtual machines in paired regions
- B. Azure Traffic Manager to route incoming traffic
- C. warm standby virtual machines in an Azure Availability Set that uses geo-redundant storage (GRS)
- D. an internal load balancer to route Internet traffic
- E. warm standby virtual machines in Azure Availability Zones

Answer: AC

Explanation:

A: An Azure Region Pair is a relationship between two Azure Regions within the same geographic region for disaster recovery purposes. If one of the regions were to experience a disaster or failure, then the services in that region will automatically failover to that regions secondary region in the pair.

C: For increased availability, you can deploy two VMs with two HANA instances within an Azure availability set that uses HANA system replication for availability.

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/sap-hana-availability-one-region>

QUESTION 25

You plan to deploy an SAP environment on Azure that will use Azure Availability Zones.

[AZ-120 Exam Dumps](#) [AZ-120 Exam Questions](#) [AZ-120 PDF Dumps](#) [AZ-120 VCE Dumps](#)

<https://www.braindump2go.com/az-120.html>

Which load balancing solution supports the deployment?

- A. Azure Basic Load Balancer
- B. Azure Standard Load Balancer
- C. Azure Application Gateway v1 SKU

Answer: B

Explanation:

When you deploy Azure VMs across Availability Zones and establish failover solutions within the same Azure region, some restrictions apply:

You can't use an Azure Basic Load Balancer to create failover cluster solutions based on Windows Server Failover Clustering or Linux Pacemaker. Instead, you need to use the Azure Standard Load Balancer SKU.

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/sap-ha-availability-zones>

QUESTION 26

You have an Azure subscription.

Your company has an SAP environment that runs on SUSE Linux Enterprise Server (SLES) servers and SAP HANA. The environment has a primary site and a disaster recovery site. Disaster recovery is based on SAP HANA system replication. The SAP ERP environment is 4 TB and has a projected growth of 5% per month.

The company has an uptime Service Level Agreement (SLA) of 99.99%, a maximum recovery time objective (RTO) of four hours, and a recovery point objective (RPO) of 10 minutes.

You plan to migrate to Azure.

You need to design an SAP landscape for the company.

Which options meet the company's requirements?

- A. Azure virtual machines and SLES for SAP application servers
SAP HANA on Azure (Large Instances) that uses SAP HANA system replication for high availability and disaster recovery
- B. ASCS/ERS and SLES clustering that uses the Pacemaker fence agent
SAP application servers deployed to an Azure Availability Zone
SAP HANA on Azure (Large Instances) that uses SAP HANA system replication for database high availability and disaster recovery
- C. SAP application instances deployed to an Azure Availability Set
SAP HANA on Azure (Large Instances) that uses SAP HANA system replication for database high availability and disaster recovery
- D. ASCS/ERS and SLES clustering that uses the Azure fence agent
SAP application servers deployed to an Azure Availability Set
SAP HANA on Azure (Large Instances) that uses SAP HANA system replication for database high availability and disaster recovery

Answer: B

Explanation:

With Availability Zones, Azure offers industry best 99.99% VM uptime SLA.

References:

<https://docs.microsoft.com/en-us/azure/traffic-manager/traffic-manager-faqs>

QUESTION 27

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

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

You deploy SAP HANA on Azure (Large Instances).

You need to back up the SAP HANA database to Azure.

Solution: You configure DB13 to back up directly to a local disk.

[AZ-120 Exam Dumps](#) [AZ-120 Exam Questions](#) [AZ-120 PDF Dumps](#) [AZ-120 VCE Dumps](#)

<https://www.braindump2go.com/az-120.html>

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

You need to back up the SAP HANA database to Azure, not to a local disk.

References:

<https://docs.microsoft.com/en-us/azure/backup/sap-hana-db-about>

<https://docs.microsoft.com/en-us/azure/backup/backup-azure-sap-hana-database#configure-backup>

QUESTION 28

You plan to deploy an SAP environment on Azure.

During a bandwidth assessment, you identify that connectivity between Azure and an on-premises datacenter requires up to 5 Gbps.

You need to identify which connectivity method you must implement to meet the bandwidth requirement.

The solution must minimize costs.

Which connectivity method should you identify?

- A. an ExpressRoute connection
- B. an Azure site-to-site VPN that is route-based
- C. an Azure site-to-site VPN that is policy-based
- D. Global VNet peering

Answer: B

Explanation:

Azure site-to-site VPN is cheaper.

Incorrect Answers:

A: ExpressRoute could be quite expensive.

C: Policy-based gateways use static routing, and only work with site-to-site connections.

References:

<https://docs.microsoft.com/en-us/azure/architecture/reference-architectures/hybrid-networking/vpn>

QUESTION 29

You plan to migrate an SAP environment to Azure.

You need to create a design to facilitate end-user access to SAP applications over the Internet, while restricting user access to the virtual machines of the SAP application servers.

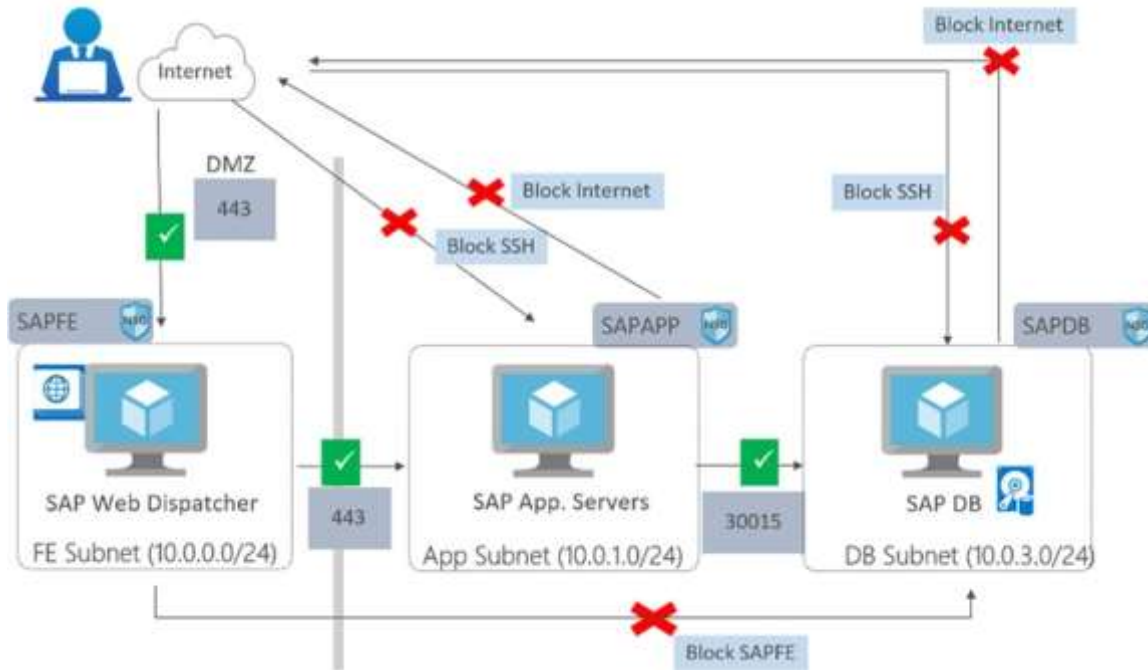
What should you include in the design?

- A. Configure a public IP address for each SAP application server
- B. Deploy an internal Azure Standard Load Balancer for incoming connections
- C. Use an SAP Web Dispatcher to route all incoming connections
- D. Configure point-to-site VPN connections for each user

Answer: C

Explanation:

1. A public internet user can reach the SAP Web-Dispatcher over port 443
2. The SAP Web-Dispatcher can reach the SAP Application server over port 443
3. The App Subnet accepts traffic on port 443 from 10.0.0.0/24
4. The SAP Application server sends traffic on port 30015 to the SAP DB server
5. The DB subnet accepts traffic on port 30015 from 10.0.1.0/24.
6. Public Internet Access is blocked on both App Subnet and DB Subnet.



References:

<https://azure.microsoft.com/en-in/blog/sap-on-azure-architecture-designing-for-security/>

QUESTION 30

You have an Azure subscription.

You deploy Active Directory domain controllers to Azure virtual machines.

You plan to deploy Azure for SAP workloads.

You plan to segregate the domain controllers from the SAP systems by using different virtual networks.

You need to recommend a solution to connect the virtual networks. The solution must minimize costs.

What should you recommend?

- A. a site-to-site VPN
- B. virtual network peering
- C. user-defined routing
- D. ExpressRoute

Answer: C

Explanation:

You can create custom, or user-defined, routes in Azure to override Azure's default system routes, or to add additional routes to a subnet's route table.

In Azure, you create a route table, then associate the route table to zero or more virtual network subnets.

Incorrect Answers:

D: ExpressRoute is a costly solution.

QUESTION 31

You deploy an SAP environment on Azure.

Your company has a Service Level Agreement (SLA) of 99.99% for SAP.

You implement Azure Availability Zones that have the following components:

- Redundant SAP application servers
- ASCS/ERS instances that use a failover cluster
- Database high availability that has a primary instance and a secondary instance

You need to validate the high availability configuration of the ASCS/ERS cluster.

What should you use?

[AZ-120 Exam Dumps](#) [AZ-120 Exam Questions](#) [AZ-120 PDF Dumps](#) [AZ-120 VCE Dumps](#)

<https://www.braindump2go.com/az-120.html>

- A. SAP Web Dispatcher
- B. Azure Traffic Manager
- C. SAPControl
- D. SAP Solution Manager

Answer: B

Explanation:

Incorrect Answers:

C: You can use SAPControl to start or stop an SAP system from the command line.

References:

<https://docs.microsoft.com/en-us/azure/architecture/reference-architectures/sap/sap-netweaver>

QUESTION 32

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: CD

Explanation:

To create a swap file in the directory that's defined by the ResourceDisk.MountPoint parameter, you can update the /etc/waagent.conf file by setting the following three parameters:

ResourceDisk.Format=y

ResourceDisk.EnableSwap=y

ResourceDisk.SwapSizeMB=xx

References:

<https://support.microsoft.com/en-us/help/4010058/how-to-add-a-swap-file-in-linux-azure-virtual-machines>

QUESTION 33

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:

To reference a credential stored in Azure Key Vault, you need to:

1. Retrieve data factory managed identity
2. Grant the managed identity access to your Azure Key Vault
3. Create a linked service pointing to your Azure Key Vault.
4. Create data store linked service, inside which reference the corresponding secret stored in key vault.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/data-factory/store-credentials-in-key-vault>

[AZ-120 Exam Dumps](#) [AZ-120 Exam Questions](#) [AZ-120 PDF Dumps](#) [AZ-120 VCE Dumps](#)

<https://www.braindump2go.com/az-120.html>