CERT EMPIRE

Microsoft AZ-305 Exam Questions

Total Questions: 300+
Demo Questions: 35
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit:
Microsoft AZ-305 Exam Questions by Cert Empire

Question: 1

You have an Azure subscription that contains a storage account. An application sometimes writes duplicate files to the storage account. You have a PowerShell script that identifies and deletes duplicate files in the storage account. Currently, the script is run manually after approval from the operations manager. You need to recommend a serverless solution that performs the following actions: Runs the script once an hour to identify whether duplicate files exist Sends an email notification to the operations manager requesting approval to delete the duplicate files Processes an email response from the operations manager specifying whether the deletion was approved Runs the script if the deletion was approved What should you include in the recommendation?

A. Azure Logic Apps and Azure Functions

B. Azure Pipelines and Azure Service Fabric

C. Azure Logic Apps and Azure Event Grid

D. Azure Functions and Azure Batch

Answer:

A

Explanation:

This scenario requires a serverless solution for both workflow orchestration and code execution. Azure Logic Apps is the ideal service for orchestrating the workflow. It can be configured with a Recurrence trigger to run hourly. Its built-in connectors for Office 365 Outlook can manage the approval email process, waiting for a response before proceeding. Azure Functions is the best choice for running the PowerShell script in a serverless, on-demand manner. The Logic App can call an Azure Function to identify the files and then, based on the manager's approval, call another Function to perform the deletion. This combination directly addresses all requirements using the most appropriate serverless tools.

Why Incorrect Options are Wrong:

B. Azure Pipelines and Azure Service Fabric: Azure Pipelines is a CI/CD tool, not a general-purpose workflow orchestrator. Service Fabric is a complex microservices platform, which is overkill for this task.
C. Azure Logic Apps and Azure Event Grid: Azure Event Grid is an event routing service. It is not needed for a time-based trigger or for executing the PowerShell script.
D. Azure Functions and Azure Batch: Azure Batch is designed for large-scale parallel and high-performance computing (HPC) workloads, not for a simple, stateful approval workflow.

References:

1. Azure Logic Apps - Workflow Orchestration: Azure Logic Apps is a cloud-based platform for creating and running automated workflows that integrate apps, data, services, and systems. The "Send approval email" action is a built-in operation that is part of the Office 365 Outlook connector.
Microsoft Learn. (2023). Create approval-based workflows by using Azure Logic Apps. Section: "Prerequisites" and "Add the Office 365 Outlook trigger".
2. Azure Functions - Serverless Code Execution: Azure Functions is a serverless compute service that lets you run event-triggered code without having to explicitly provision or manage infrastructure. It natively supports PowerShell, making it suitable for executing the script.
Microsoft Learn. (2023). An introduction to Azure Functions. Section: "What is Functions?".
3. Integration of Logic Apps and Functions: Logic Apps can call Azure Functions to add custom code to a workflow. This is a standard design pattern for extending Logic Apps capabilities.
Microsoft Learn. (2023). Call Azure Functions from Azure Logic Apps. Section: "Create a function in the Azure portal".
4. Scheduled Triggers: Both Logic Apps (Recurrence trigger) and Functions (Timer trigger) support scheduled execution, fulfilling the "once an hour" requirement.
Microsoft Learn. (2023). Run tasks and workflows on a schedule with Azure Logic Apps. Section: "Add the Recurrence trigger".

Question: 2

DRAG DROP You have an on-premises network that uses on IP address space of 172.16.0.0/16 You plan to deploy 25 virtual machines to a new azure subscription. You identity the following technical requirements. All Azure virtual machines must be placed on the same subnet subnet1. All the Azure virtual machines must be able to communicate with all on premises severs. The servers must be able to communicate between the on-premises network and Azure by using a site to site VPN. You need to recommend a subnet design that meets the technical requirements. What should you include in the recommendation? To answer, drag the appropriate network addresses to the correct subnet. Each network address 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.

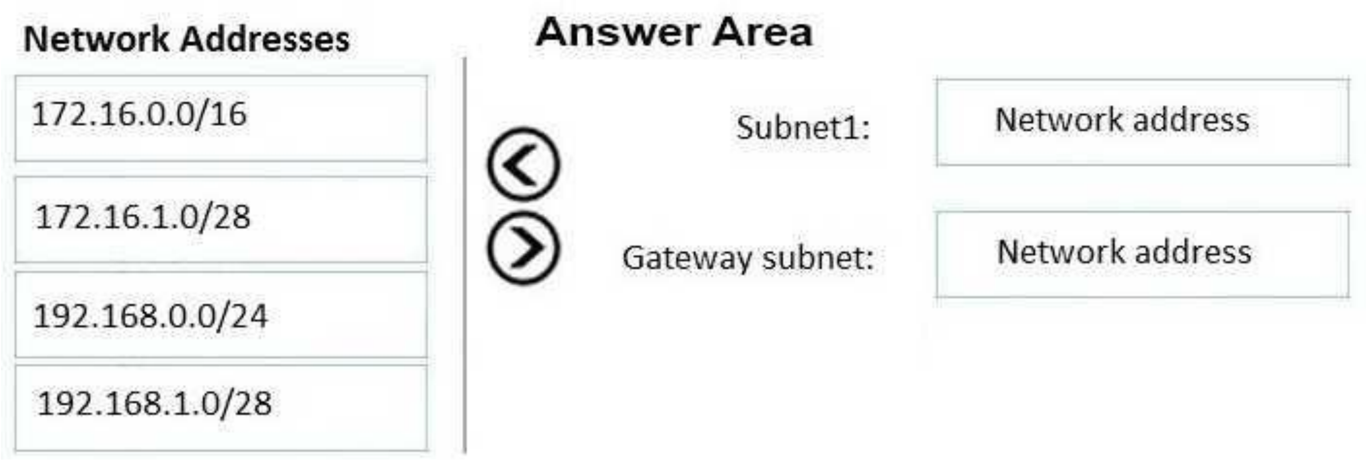

Network Addresses

172.16.0.0/16

172.16.1.0/28

192.168.0.0/24

192.168.1.0/28

Answer Area

Subnet1: Network address

Gateway subnet: Network address

Answer:

Subnet1: 192.168.0.0/24

Gateway subnet: 192.168.1.0/28

Explanation:

The solution is based on two key Azure networking principles: non-overlapping address spaces for VPN connectivity and subnet sizing.

- Non-Overlapping Address Space: The on-premises network uses the 172.16.0.0/16 address space. To establish a Site-to-Site VPN connection and ensure proper routing, the Azure virtual network's address space must not overlap with the on-premises space. This requirement immediately eliminates 172.16.0.0/16 and 172.16.1.0/28 as options for any Azure subnet.

- Subnet Sizing:

- Subnet1: This subnet must host 25 virtual machines. Azure reserves 5 IP addresses in every subnet for protocol conformance. Therefore, a subnet must contain at least 25 + 5 = 30 total IP addresses.

- The 192.168.1.0/28 subnet provides only 2(3228)=16 total IPs, which is insufficient.

- The 192.168.0.0/24 subnet provides 2(3224)=256 total IPs (251 usable for VMs), which is sufficient.

- Gateway Subnet: A Site-to-Site VPN requires a dedicated subnet for the virtual network gateway, conventionally named GatewaySubnet. From the remaining valid options, 192.168.1.0/28 is the only logical choice for this purpose.

References:

Microsoft Learn, "Plan virtual networks":
In the "IP addresses" section, the documentation states, "When you create a VNet, you must specify a custom private IP address space... The address spaces you assign to VNet and on-premises can't overlap. If they overlap, connectivity between them isn't possible." This supports the decision to discard the 172.16.x.x addresses.
Microsoft Learn, "Virtual Network frequently asked questions (FAQ)":

Under the "How many IP addresses are reserved within a subnet?" section, it is specified that "Azure reserves some IP addresses within each subnet. The first four addresses and the last address are reserved, for a total of 5 reserved IP addresses." This confirms the calculation used for determining the required size of Subnet1.
Microsoft Learn, "About VPN Gateway configuration settings":
The "GatewaySubnet" section explains the need for a dedicated subnet for the virtual network gateway. It states, "The gateway subnet contains the IP addresses that the virtual network gateway services use... you need to create a gateway subnet for your virtual network." This justifies the allocation of a separate network address for the Gateway subnet.

Question: 3

You are designing an Azure solution. The network traffic for the solution must be securely distributed by providing the following features: HTTPS protocol Round robin routing SSL offloading You need to recommend a load balancing option. What should you recommend?

A. Azure Load Balancer

B. Azure Traffic Manager

C. Azure Internal Load Balancer (ILB)

D. Azure Application Gateway

Answer:

D

Explanation:

Azure Application Gateway is a web traffic load balancer that operates at the application layer (Layer 7). It is specifically designed to manage HTTP and HTTPS traffic. It fully supports the required features: it handles the HTTPS protocol, can terminate the SSL/TLS connection at the gateway (SSL offloading) to reduce the load on backend web servers, and distributes traffic to backend pool members using a round-robin mechanism. The other options operate at different layers or serve different purposes and cannot meet all the specified requirements.

Why Incorrect Options are Wrong:

A. Azure Load Balancer: This is a Layer 4 (Transport Layer) service that distributes TCP/UDP traffic. It does not have the capability to perform SSL offloading as it is unaware of the application-level data.

B. Azure Traffic Manager: This is a DNS-based traffic load balancer. It routes traffic at the DNS level to different endpoints but does not process or terminate the actual traffic, and therefore cannot perform SSL offloading.

C. Azure Internal Load Balancer (ILB): This is a variant of the standard Azure Load Balancer (Layer 4) used for private, internal networks. It shares the same limitation of not being able to perform application-level tasks like SSL offloading.

References:

1. Microsoft Docs, Azure Application Gateway. "What is Azure Application Gateway?". Under the "Features" section, it lists "SSL/TLS termination" and "Round-robin routing". It states, "Application Gateway is a web traffic load balancer that enables you to manage traffic to your web applications."

2. Microsoft Docs, Azure Application Gateway. "TLS termination and end to end TLS with Application Gateway". In the "Overview" section, it explicitly states, "Application Gateway

supports TLS/SSL termination at the gateway, after which traffic typically flows unencrypted to the backend servers."

3. Microsoft Docs, Azure Load Balancer. "What is Azure Load Balancer?". The "Overview" section clearly defines it: "Azure Load Balancer is a high-performance, ultra low-latency Layer 4 load-balancing service (inbound and outbound) for all UDP and TCP protocols."

4. Microsoft Docs, Azure Architecture Center. "Choose a load-balancing service". The comparison table under the "Summary of services" section shows that only Application Gateway supports "SSL/TLS offloading" and is a "Layer 7 (HTTP/S)" service, whereas Azure Load Balancer is Layer 4 and Traffic Manager operates at the DNS layer.

Question: 4

Your company, named Contoso, Ltd, implements several Azure logic apps that have HTTP triggers: The logic apps provide access to an on-premises web service. Contoso establishes a partnership with another company named Fabrikam, Inc. Fabrikam does not have an existing Azure Active Directory (Azure AD) tenant and uses third-party OAuth 2.0 identity management to authenticate its users. Developers at Fabrikam plan to use a subset of the logics apps to build applications that will integrate with the on-premises web service of Contoso. You need to design a solution to provide the Fabrikam developers with access to the logic apps. The solution must meet the following requirements: Requests to the logic apps from the developers must be limited to lower rates than the requests from the users at Contoso. The developers must be able to rely on their existing OAuth 2.0 provider to gain access to the logic apps. The solution must NOT require changes to the logic apps. The solution must NOT use Azure AD guest accounts. What should you include in the solution?

A. Azure AD business-to-business (B2B)

B. Azure Front Door

C. Azure API Management

D. Azure AD Application Proxy

Answer:

C

Explanation:

Azure API Management (APIM) is designed to act as a secure gateway for backend services, such as the Azure Logic Apps in this scenario. APIM can be placed in front of the Logic Apps without requiring any changes to them. It allows for the creation of policies to manage access, including rate-limiting policies that can be applied specifically to Fabrikam's developers. Furthermore, APIM's validate-jwt policy can be configured to validate tokens from Fabrikam's external OAuth 2.0 provider by specifying the provider's OpenID Connect metadata endpoint. This meets the authentication requirement without creating Azure AD guest accounts, fulfilling all constraints of the question.

Why Incorrect Options are Wrong:

A. Azure AD business-to-business (B2B) is incorrect because it functions by inviting external users as guests into the Azure AD tenant, which is explicitly prohibited by the requirements.
B. Azure Front Door is less suitable because its primary functions are global load balancing and WAF protection, not the granular, identity-based API policy enforcement and third-party identity federation required here.
D. Azure AD Application Proxy is incorrect as its purpose is to publish on-premises applications

for secure remote access, not to manage and secure cloud-based APIs for external partners.

References:

1. Azure API Management Policies:
Rate Limiting: The rate-limit-by-key policy can be used to limit the call rate for Fabrikam's developers based on a subscription key.
Source: Microsoft Docs, "Azure API Management access restriction policies", Section: "Rate limit by key".
Third-Party OAuth 2.0 Integration: The validate-jwt policy is used to enforce the validity of a JWT from an external identity provider.
Source: Microsoft Docs, "Azure API Management access restriction policies", Section: "Validate JWT". The documentation states, "Use this policy to enforce the existence and validity of a JSON Web Token (JWT)..." and shows configuration using an openid-config URL, which is standard for any OpenID Connect compliant provider.
2. Azure AD B2B: This service explicitly uses guest accounts.
Source: Microsoft Docs, "What is Azure AD B2B collaboration?". The overview states, "Azure Active Directory (Azure AD) B2B collaboration is a feature... that lets you invite guest users to collaborate in your organization."
3. Azure AD Application Proxy: This service is for on-premises applications.
Source: Microsoft Docs, "Remote access to on-premises applications through Azure AD Application Proxy". The article summary clearly states its purpose is to provide "secure remote access to on-premises web applications."

Question: 5

You need to design a solution that will execute custom C# code in response to an event routed to Azure Event Grid. The solution must meet the following requirements: The executed code must be able to access the private IP address of a Microsoft SQL Server instance that runs on an Azure virtual machine. Costs must be minimized. What should you include in the solution?

A. Azure Logic Apps in the integrated service environment

B. Azure Functions in the Dedicated plan and the Basic Azure App Service plan

C. Azure Logic Apps in the Consumption plan

D. Azure Functions in the Consumption plan

Answer:

B

Explanation:

The solution requires executing custom C# code, for which Azure Functions is the most suitable service. A critical requirement is accessing a private IP address of a SQL Server on a virtual machine, which necessitates integrating the execution environment with an Azure Virtual Network (VNet).

Azure Functions running on a Consumption plan do not support VNet integration for outbound traffic. This capability is only available in the Premium and Dedicated (App Service) plans. While an Integrated Service Environment (ISE) for Logic Apps also provides VNet connectivity, it is a significantly more expensive, single-tenant solution. Therefore, the Dedicated plan for Azure Functions is the lowest-cost option presented that meets all the mandatory technical requirements.

Why Incorrect Options are Wrong:

A. Azure Logic Apps in the integrated service environment: This option meets the VNet requirement but is a high-cost, premium offering that violates the principle of cost minimization.

C. Azure Logic Apps in the Consumption plan: This plan does not support direct execution of custom C# code and, more importantly, lacks the VNet integration required to access private IP addresses.

D. Azure Functions in the Consumption plan: This is the most cost-effective plan for event-driven code, but it does not support VNet integration and therefore cannot access resources by their private IP address within a VNet.

References:

1. Azure Functions networking options: The official documentation explicitly states which plans support VNet integration.
Microsoft Docs. (2023). Azure Functions networking options. "Virtual network integration". The documentation states, "The virtual network integration feature is available in Azure Functions for apps in a Premium plan, or in an App Service plan (Standard or higher)." This confirms that the Consumption plan (Option D) cannot fulfill the requirement.
Reference: https://learn.microsoft.com/en-us/azure/azure-functions/functions-networking-options#virtual-network-integration
2. Azure App Service plan overview: This document details the capabilities of different App Service plans, which host Azure Functions on the Dedicated plan.
Microsoft Docs. (2023). Azure App Service plan overview. "App Service plan". This page outlines that plans like Basic, Standard, and Premium support features such as VNet integration, differentiating them from the serverless Consumption model.
Reference: https://learn.microsoft.com/en-us/azure/app-service/overview-hosting-plans
3. Connect to Azure virtual networks from Azure Logic Apps: This document clarifies networking for Logic Apps.
Microsoft Docs. (2023). Connect to Azure virtual networks from Azure Logic Apps. "Access virtual network resources". The documentation specifies that to access resources inside a VNet, you need a single-tenant Logic App (Standard plan) or an Integration Service Environment (ISE), not the multi-tenant Consumption plan (Option C). The ISE (Option A) is described as a premium, isolated environment.
Reference: https://learn.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-single-tenant

Question: 6

The developers at your company are building a containerized Python Django app. You need to recommend platform to host the app. The solution must meet the following requirements: Support autoscaling. Support continuous deployment from an Azure Container Registry. Provide built-in functionality to authenticate app users by using Azure Active Directory (Azure AD). Which platform should you include in the recommendation?

A. Azure Container instances

B. an Azure App Service instance that uses containers

C. Azure Kubernetes Service (AKS)

Answer:

B

Explanation:

Azure App Service for Containers is the most suitable platform as it directly meets all the specified requirements. It provides robust, built-in autoscaling capabilities for the App Service Plan. It features a native continuous deployment option that can be configured to automatically pull and deploy new container images from an Azure Container Registry (ACR). Most importantly, App Service offers a built-in "Authentication / Authorization" feature (often called "Easy Auth") that allows developers to secure the application with Azure Active Directory (Azure AD) through simple configuration, without requiring any code changes in the app itself.

Why Incorrect Options are Wrong:

A. Azure Container instances: This service is designed for single, isolated containers and lacks built-in autoscaling and application-level authentication features required by the scenario.

C. Azure Kubernetes Service (AKS): While AKS supports autoscaling and ACR integration, it does not provide built-in platform functionality to authenticate app users. Securing the app would require manual implementation within the application code or complex configuration of an ingress controller and authentication proxies.

References:

1. Azure App Service Authentication: Microsoft Docs, "Authentication and authorization in Azure App Service and Azure Functions". Under the "How it works" section, it states, "The authentication and authorization module runs in the same sandbox as your application code. When it's enabled, every incoming HTTP request passes through it before being handled by your application code." This confirms the built-in, codeless nature of the feature.
2. Azure App Service Continuous Deployment from ACR: Microsoft Docs, "Deploy from Azure Container Registry". The document outlines the steps to "Configure a webhook in your web app

that will receive push notifications from your registry," enabling continuous deployment.
3. Azure App Service Autoscaling: Microsoft Docs, "Scale instance count manually or automatically". This document details how to configure autoscale rules for an App Service plan based on metrics like CPU percentage, which directly addresses the autoscaling requirement.
4. Azure Kubernetes Service (AKS) Azure AD Integration: Microsoft Docs, "Use Azure Active Directory in Azure Kubernetes Service". The introduction explicitly states, "AKS can be configured to use Azure AD for user authentication. In this configuration, you can sign in to an AKS cluster by using an Azure AD authentication token... This document shows you how to create the necessary Azure AD components, then deploy an Azure AD-enabled cluster and sign in to the AKS cluster." This clarifies that AAD integration is for authenticating to the Kubernetes API, not for the applications running within the cluster.

Question: 7

You have an on-premises network to which you deploy a virtual appliance. You plan to deploy several Azure virtual machines and connect the on-premises network to Azure by using a Site-to-Site connection. All network traffic that will be directed from the Azure virtual machines to a specific subnet must flow through the virtual appliance. You need to recommend solutions to manage network traffic. Which two options should you recommend? Each correct answer presents a complete solution.

A. Configure Azure Traffic Manager.

B. Implement an Azure virtual network.

C. Implement Azure ExpressRoute.

D. Configure a routing table.

Answer:

C, D

Explanation:

The objective is to direct traffic from Azure virtual machines to a specific on-premises subnet through an on-premises virtual appliance. This requires overriding Azure's default routing behavior.

1. Configure a routing table (D): This is a standard method for controlling traffic flow in Azure. By creating a Route Table with a User-Defined Route (UDR), you can specify that traffic destined for the on-premises subnet prefix must be sent to the VirtualNetworkGateway as the next hop. This forces the traffic across the Site-to-Site connection, where on-premises routing can then direct it through the virtual appliance.
2. Implement Azure ExpressRoute (C): ExpressRoute provides a private, dedicated connection that uses the Border Gateway Protocol (BGP) for dynamic route exchange. Using BGP, you can advertise specific routes from your on-premises network to Azure. This allows your on-premises network to announce the route to the target subnet, influencing Azure's routing decisions to ensure traffic is sent over the connection and subsequently through your on-premises appliance.

Why Incorrect Options are Wrong:

A. Configure Azure Traffic Manager: This is a DNS-based load balancer that operates at the application layer to direct client requests to endpoints. It does not control network-level (IP) routing.

B. Implement an Azure virtual network: An Azure virtual network is a fundamental prerequisite for deploying the virtual machines, but it does not, by itself, provide the specific routing controls required.

References:

1. Microsoft Documentation, "Virtual network traffic routing": This document explains how Azure routes traffic and how User-Defined Routes (UDRs) in a route table can override default behaviors. It states, "You can create custom, or user-defined(static), routes in Azure to override Azure's default system routes... You can create a route table with a route that sends traffic destined for an on-premises network to your virtual network gateway." This supports option D.
2. Microsoft Documentation, "About BGP with Azure VPN Gateway": This document details the use of BGP for dynamic routing over hybrid connections. It notes, "When you use BGP with an Azure VPN Gateway or ExpressRoute, you advertise your on-premises prefixes to Microsoft Azure." This allows for granular control over traffic paths, supporting the solution described in option C.
3. Microsoft Documentation, "ExpressRoute routing requirements": This document confirms that ExpressRoute uses BGP to "exchange routes between your on-premises network, your instances in Azure, and Microsoft public addresses." This is the mechanism that enables the traffic management solution.

Question: 8

You are designing a message application that will run on an on-premises Ubuntu virtual machine. The application will use Azure Storage queues. You need to recommend a processing solution for the application to interact with the storage queues. The solution must meet the following requirements: Create and delete queues daily. Be scheduled by using a CRON job. Upload messages every five minutes. What should developers use to interact with the queues?

A. Azure CLI

B. AzCopy

C. Azure Data Factory

D. .NET Core

Answer:

D

Explanation:

The question requires a solution for an "application" that developers will build. For developing an application that programmatically interacts with Azure services, the standard and recommended approach is to use an Azure Software Development Kit (SDK). .NET Core is a cross-platform framework that runs on Ubuntu and has a robust Azure Storage SDK. This allows developers to integrate queue operations (create, delete, send messages) directly into their application code, providing granular control, complex logic, and sophisticated error handling. The resulting compiled application can then be scheduled to run using a CRON job, fulfilling all the requirements in a maintainable and scalable manner.

Why Incorrect Options are Wrong:

A. Azure CLI: While the Azure CLI can perform the required operations and be scripted, it is primarily a tool for interactive management and automation, not for building the core logic of a production application. Using an SDK is the standard for application development.
B. AzCopy: AzCopy is a specialized command-line utility designed for high-performance bulk copying of data to and from Azure Blob and File storage. It does not support queue management or message operations.
C. Azure Data Factory: This is a cloud-based data orchestration service for building ETL/ELT pipelines. It is not a tool that runs on an on-premises virtual machine or is invoked directly by a local CRON job.

References:

1. Azure SDKs vs. CLI: Microsoft's official documentation, "Choose the right Azure command-line tool," clarifies the use cases. It states, "The Azure SDKs are collections of libraries that developers use to programmatically interact with Azure services... The Azure CLI is for the automation of administrative tasks, most often in the context of a CI/CD pipeline." This supports using an SDK (.NET Core) for application development.
Source: Microsoft Docs, "Choose the right Azure command-line tool," Section: "What are the Azure SDKs?".
2. .NET Core SDK for Queues: The quickstart for the Azure Queue Storage client library for .NET demonstrates the standard programmatic approach for an application to create queues and send messages, which aligns with the requirements.
Source: Microsoft Docs, "Quickstart: Azure Queue Storage client library for .NET," Section: "Code examples".
3. AzCopy Purpose: The official documentation for AzCopy explicitly states its function: "You can use AzCopy to copy blobs or files to or from a storage account." It does not mention queue operations.
Source: Microsoft Docs, "Get started with AzCopy," Introduction paragraph.

Question: 9

HOTSPOT You plan to deploy the backup policy shown in the following exhibit.

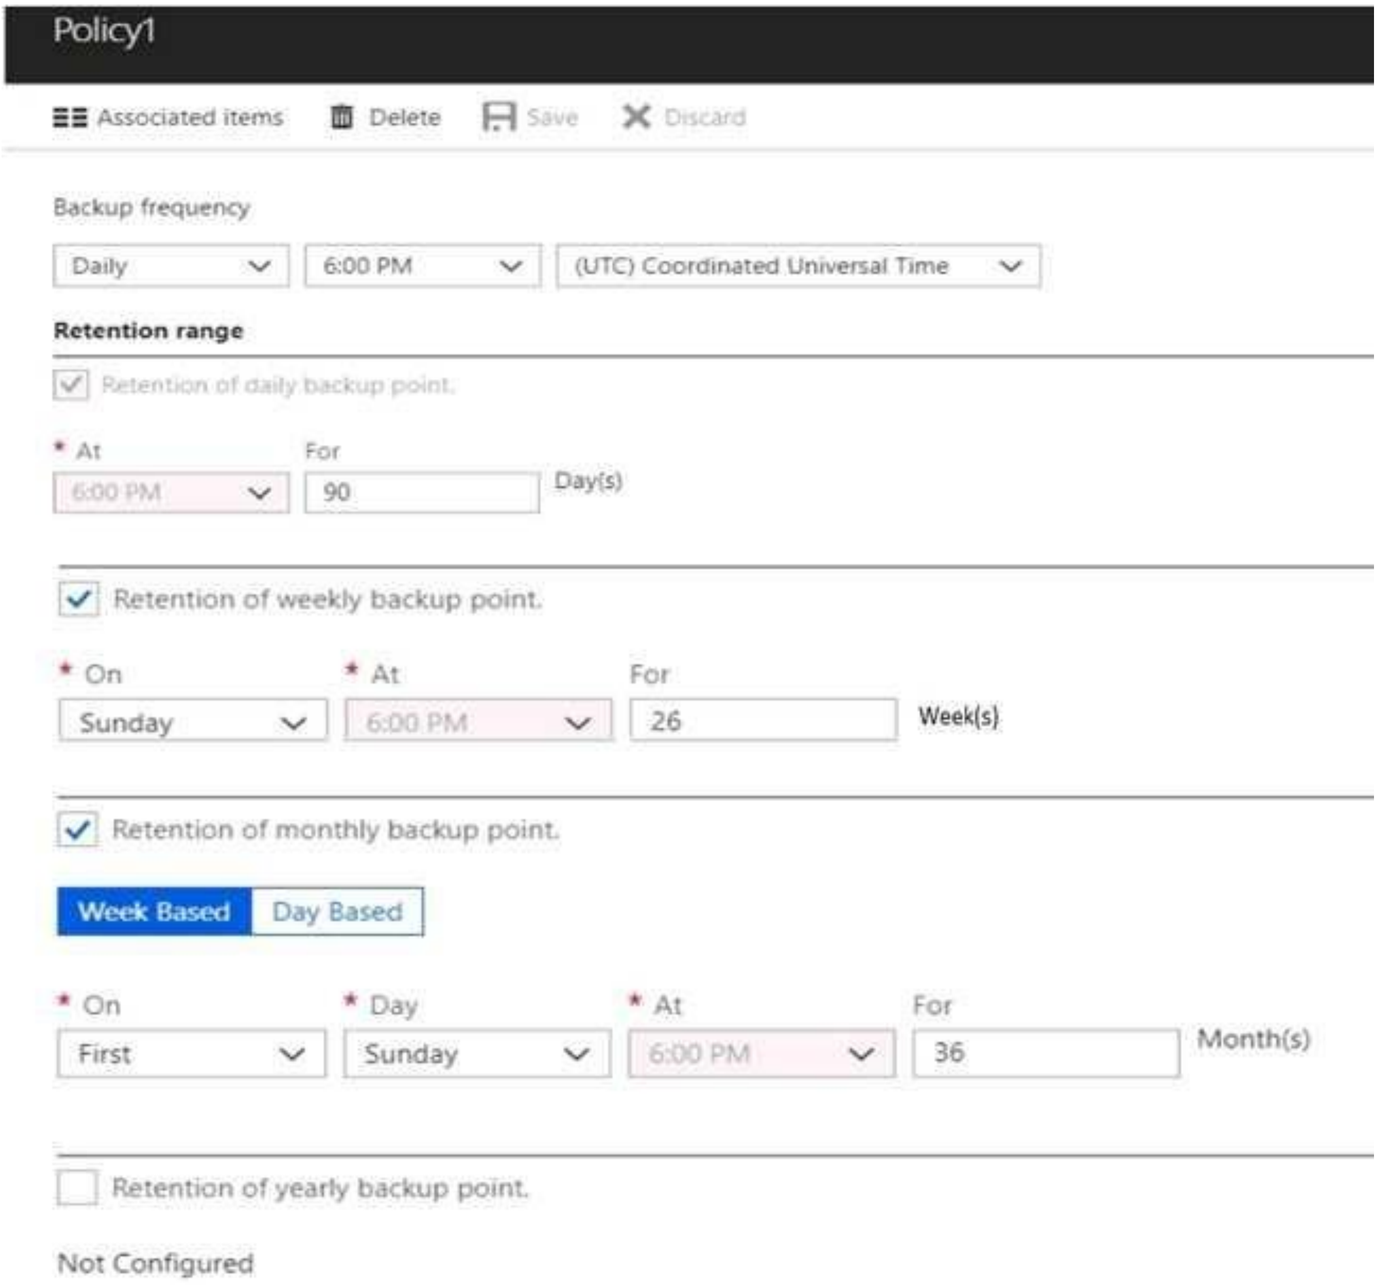

Policy1

Backup frequency

Daily | 6:00 PM | (UTC) Coordinated Universal Time

Retention range

☑ Retention of daily backup point.

* At: 6:00 PM — For: 90 Day(s)

☑ Retention of weekly backup point.

* On: Sunday — * At: 6:00 PM — For: 26 Week(s)

☑ Retention of monthly backup point.

* On: First — * Day: Sunday — * At: 6:00 PM — For: 36 Month(s)

□ Retention of yearly backup point.

Not Configured

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

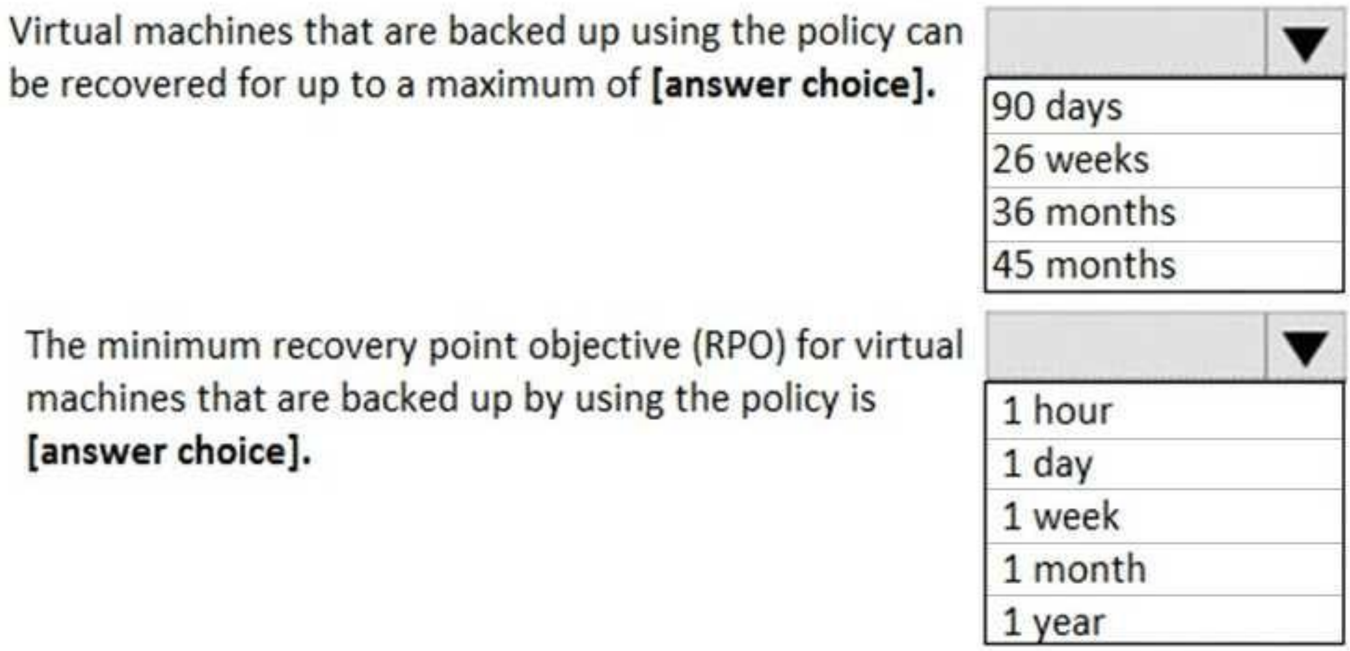

Virtual machines that are backed up using the policy can be recovered for up to a maximum of **[answer choice].**

- 90 days
- 26 weeks
- 36 months
- 45 months

The minimum recovery point objective (RPO) for virtual machines that are backed up by using the policy is **[answer choice].**

- 1 hour
- 1 day
- 1 week
- 1 month
- 1 year

Explanation:

Hotspot 1

Correct Answer: 10 years

Explanation: The backup taken on the first Sunday of January qualifies for all four retention rules: daily (180 days), weekly (104 weeks), monthly (60 months), and yearly (10 years). Azure Backup applies a Grandfather-Father-Son (GFS) retention logic. When a recovery point is marked for retention by multiple rules, it is stored for the longest duration specified among those rules. In this case, the 10-year retention period is the longest and will be applied to this specific backup.

Hotspot 2

Correct Answer: 104 weeks

Explanation: The backup taken on the second Sunday of January qualifies for the daily (180 days) and weekly (104 weeks) retention rules. It does not meet the criteria for the monthly rule ("First Sunday") or the yearly rule ("First Sunday of January"). Following the GFS principle, Azure Backup applies the longest applicable retention period. Between 180 days and 104 weeks, the 104-week period is longer. Therefore, this backup will be retained for 104 weeks.

References

1. Microsoft Learn. (2023). Back up an Azure VM from the VM settings. "How does retention work?" section. Retrieved from https://learn.microsoft.com/en-us/azure/backup/backup-azure-vms

-first-look-arm#how-does-retention-work.
Quote: "When a recovery point is tagged for retention for more than one retention rule, it's stored for the longest duration." This directly confirms the principle that the longest retention period is applied when multiple rules match.

2. Microsoft Learn. (2024). Create and manage backup policies for Azure VM backup. "Create a backup policy" section. Retrieved from https://learn.microsoft.com/en-us/azure/backup/backup-azure-manage-vms#create-a-backup-policy.
Reference: This document details the configuration of daily, weekly, monthly, and yearly retention points, which form the basis of the GFS scheme evaluated in the question. The structure of the policy creation UI shown in the documentation aligns with the policy in the exhibit.

Question: 10

DRAG DROP Your company identifies the following business continuity and disaster recovery objectives for virtual machines that host sales, finance, and reporting application in the company's on-premises data center. •The finance application requires that data be retained for seven years. In the event of a disaster, the application must be able to run from Azure. The recovery in objective (RTO) is 10 minutes, • The reporting application must be able to recover point in-time data al a daily granularity. The RTO is eight hours. •The sales application must be able to fail over to second on-premises data center. You need to recommend which Azure services meet the business community and disaster recovery objectives. The solution must minimize costs. What should you recommend for each application? To answer, drag the appropriate services to the correct application. Each service may be used owe. More than once not at an You may need to drag the spin bar between panes or scroll 10 view content.

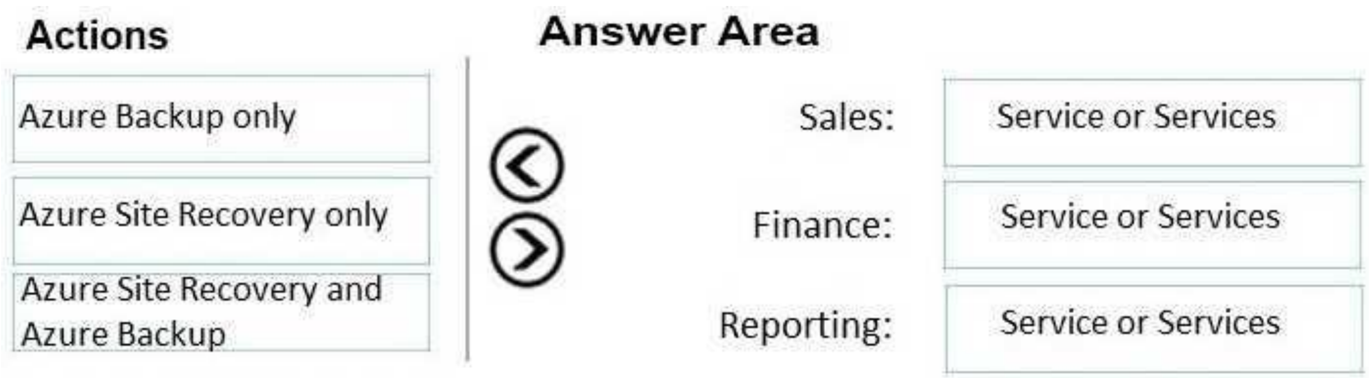

Actions

Azure Backup only

Azure Site Recovery only

Azure Site Recovery and Azure Backup

Answer Area

Sales: Service or Services

Finance: Service or Services

Reporting: Service or Services

Answer:

Sales: Azure Site Recovery only

Finance: Azure Site Recovery and Azure Backup

Reporting: Azure Backup only

Explanation:

The selection of services is based on the specific recovery and retention requirements for each application while minimizing costs.

- Sales: The requirement is to fail over to a second on-premises data center. Azure Site Recovery is the appropriate service as it can orchestrate replication and failover between two on-premises sites (e.g., Hyper-V to Hyper-V or VMware to VMware), using Azure only as the control plane. This meets the requirement without needing Azure Backup.

- Finance: This application has two distinct needs: a very low Recovery Time Objective (RTO of

10 minutes) for disaster recovery and seven-year data retention. Azure Site Recovery is necessary to meet the low RTO by replicating the VM to Azure for a quick failover. Azure Backup is required to meet the long-term data retention policy of seven years.

• Reporting: The requirements are a daily recovery point and a high RTO of eight hours. This lengthy RTO does not necessitate an immediate failover solution. Azure Backup is sufficient and most cost-effective, as it can perform daily backups and restore a full VM within the eight-hour window.

References:

Azure Site Recovery Documentation: Microsoft's official documentation outlines the scenarios supported by Azure Site Recovery. It explicitly states it supports "Replication of on-premises VMs to a secondary on-premises datacenter." This directly addresses the requirement for the Sales application.
Source: Microsoft Learn, "About Site Recovery," Scenarios section.
Azure Site Recovery & Azure Backup Synergy: The documentation clarifies the complementary roles of the two services. "Site Recovery provides disaster recovery... Azure Backup provides backup." It explains that Site Recovery contributes to your business continuity and disaster recovery (BCDR) strategy by keeping workloads available, while Backup protects data and provides long-term retention. This justifies using both for the Finance application.

Source: Microsoft Learn, "Azure business continuity technical guidance," Page 4, "Azure Site Recovery and Azure Backup."
Azure Backup Retention Policies: The Azure Backup documentation confirms its capability for long-term retention. Policies can be configured to retain data for up to 99 years, easily satisfying the seven-year requirement for the Finance application.
Source: Microsoft Learn, "What is the Azure Backup service?," "Backup and restore" section.
Azure Backup for VM Recovery: Restoring a VM from Azure Backup is a suitable solution for higher RTOs. While the exact time can vary, an RTO of eight hours, as required by the Reporting application, is well within the typical capabilities of a VM restore from a backup vault, making it the most cost-effective choice.
Source: Microsoft Learn, "About Azure VM backup," "Backup and restore process" section.

Question: 11

You need to design a highly available Azure SQL database that meets the following requirements: * Failover between replicas of the database must occur without any data loss. * The database must remain available in the event of a zone outage. * Costs must be minimized. Which deployment option should you use?

A. Azure SQL Database Business Critical

B. Azure SQL Database Managed Instance Business Critical

C. Azure SQL Database Serverless

D. Azure SQL Database Premium

Answer:

D

Explanation:

The solution requires a database that supports zone-redundant, synchronous replication for high availability with no data loss (RPO of 0), while minimizing costs. The Azure SQL Database Premium service tier, when configured with zone redundancy, meets these criteria. It uses an architecture based on Always On availability groups, which maintains multiple synchronous replicas. This model allows for fast failover between replicas within or across availability zones without data loss. While the Business Critical tier also offers this capability, the Premium tier (in the DTU model) or the lower tiers of Business Critical (in the vCore model) are generally more cost-effective, satisfying the cost minimization requirement.

Why Incorrect Options are Wrong:

A. Azure SQL Database Business Critical: This tier meets the availability and data loss requirements but is generally more expensive than the Premium tier, thus not meeting the cost minimization requirement.
B. Azure SQL Database Managed Instance Business Critical: This is a platform-as-a-service (PaaS) offering for migrating SQL Server workloads. It is typically more expensive than a single Azure SQL Database and is not the most cost-effective choice for this scenario.
C. Azure SQL Database Serverless: While the Serverless compute tier (within the General Purpose service tier) can be configured for zone redundancy with zero RPO, its high-availability model relies on redundant storage (ZRS) and standby compute capacity, not a failover "between replicas" in the same architectural sense as the Premium/Business Critical tiers.

References:

1. Microsoft Documentation - High availability for Azure SQL Database and SQL Managed Instance: This document details the architectures. For the Premium/Business Critical tiers, it states, "The premium and business critical service tiers are based on a model that integrates compute and storage on a single node... This model relies on a quorum of database engine nodes, using a technology similar to SQL Server Always On availability groups." This confirms the "failover between replicas" model. For General Purpose, it describes the separation of compute and storage.
Source: Microsoft Docs, "High availability for Azure SQL Database and SQL Managed Instance", Section: "Premium and Business Critical service tier availability".
2. Microsoft Documentation - DTU-based purchasing model overview: This document describes the Premium tier and its suitability for high-performance workloads requiring high availability. It confirms that the Premium tier is designed for I/O-intensive workloads with the highest levels of availability.
Source: Microsoft Docs, "DTU-based purchasing model overview", Section: "Service tiers".
3. Microsoft Documentation - Configure zone-redundant high availability: This document confirms that zone redundancy is available for the Premium service tier. It states, "When you provision a database or an elastic pool, you can specify that it's zone redundant. Zone redundancy is available for databases in the General Purpose, Premium, Business Critical, and Hyperscale (preview) service tiers."
Source: Microsoft Docs, "Configure zone-redundant high availability for Azure SQL Database".

Question: 12

HOTSPOT Your company deploys an Azure App Service Web App. During testing the application fails under load. The application cannot handle more than 100 concurrent user sessions. You enable the Always On feature. You also configure auto-scaling to increase counts from two to 10 based on HTTP queue length. You need to improve the performance of the application. Which solution should you use for each application scenario? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

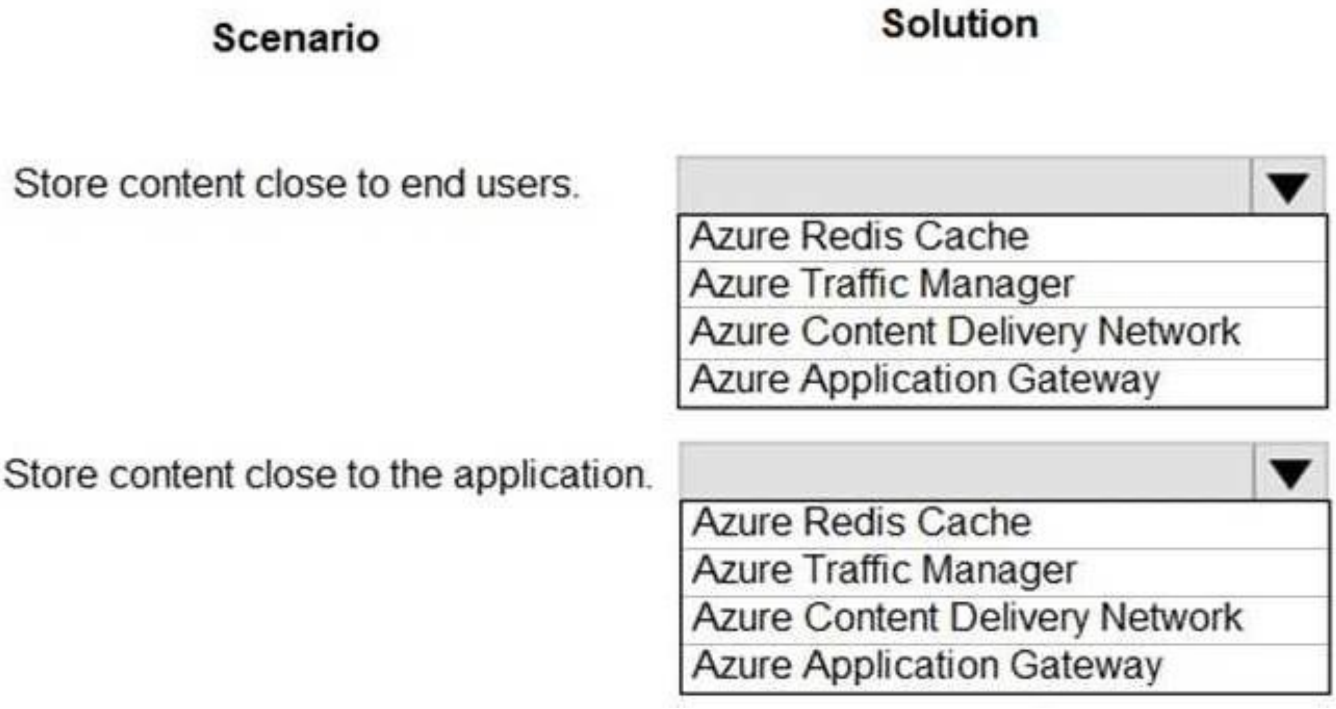

Scenario	Solution
Store content close to end users.	▼ Azure Redis Cache Azure Traffic Manager Azure Content Delivery Network Azure Application Gateway
Store content close to the application.	▼ Azure Redis Cache Azure Traffic Manager Azure Content Delivery Network Azure Application Gateway

Answer:

Store content close to end users: Azure Content Delivery Network

Store content close to the application: Azure Redis Cache

Explanation:

Azure Content Delivery Network (CDN) is the correct choice for storing content close to end users. A CDN is a globally distributed network of servers (Points of Presence or POPs) that caches static content like images, CSS, and JavaScript files. When a user requests this content, it is delivered from the geographically closest POP, which significantly reduces latency and improves load times.

Azure Cache for Redis is the appropriate solution for storing content close to the application. It provides a managed, in-memory data store that applications can use as a high-performance cache. By caching frequently accessed data in Redis, which is located in the same Azure region as the application, the application can reduce its dependency on slower backend datastores (like SQL databases), thereby improving response time and scalability.

References:

Azure Content Delivery Network (CDN) Documentation: "An Azure content delivery network (CDN) is a distributed network of servers that can efficiently deliver web content to users. CDNs store cached content on edge servers in point-of-presence (POP) locations that are close to end users, to minimize latency."
Source: Microsoft Corporation. (2023). What is a content delivery network (CDN) on Azure?. Microsoft Learn. Retrieved from https://learn.microsoft.com/en-us/azure/cdn/cdn-overview, Overview section.
Azure Cache for Redis Documentation: "Azure Cache for Redis provides an in-memory data store based on the Redis software….. It improves the performance and scalability of an application that uses backend data stores heavily. It's able to process large volumes of application requests by keeping frequently accessed data in the server memory that can be written to and read from quickly."
Source: Microsoft Corporation. (2023). What is Azure Cache for Redis?. Microsoft Learn. Retrieved from https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-overview, Introduction section.

Question: 13

You plan to store data in Azure Blob storage for many years. The stored data will be accessed rarely. You need to ensure that the data in Blob storage is always available for immediate access. The solution must minimize storage costs. Which storage tier should you use?

A. Cool

B. Archive

C. Hot

Answer:

A

Explanation:

The Cool access tier is the optimal choice for this scenario. It is specifically designed for data that is infrequently accessed but must be available for immediate retrieval. It offers a balance between lower storage costs compared to the Hot tier and instant, online access, which the Archive tier cannot provide. Since the requirements are long-term storage, rare access, immediate availability, and minimized cost, the Cool tier is the only option that satisfies all these constraints. The Archive tier fails the "immediate access" requirement, and the Hot tier fails the "minimize storage costs" requirement for rarely accessed data.

Why Incorrect Options are Wrong:

B. Archive: This is an offline tier. Data must be rehydrated before access, a process that can take several hours, thus failing the "immediate access" requirement.

C. Hot: This tier is optimized for frequently accessed data and has the highest storage costs, which contradicts the requirement to minimize costs for rarely accessed data.

References:

1. Microsoft Learn, Azure Blob storage access tiers - Hot, Cool, and Archive. In the "Summary of access tier options" table, it clearly states that both Hot and Cool tiers are "Online" with "Milliseconds" access latency. In contrast, the Archive tier is "Offline" with "Hours" of access latency (rehydration time). This document confirms Cool tier provides immediate access while being cost-effective for infrequent use.
Reference: learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#summary-of-access-tier-options

2. Microsoft Learn, Hot, Cool, and Archive access tiers for blob data. This document describes the usage patterns for each tier. It specifies, "Cool access tier - Optimized for storing large amounts of data that is infrequently accessed and stored for at least 30 days." and "Archive access tier - Optimized for storing data that is rarely accessed... with flexible latency requirements

(on the order of hours)." This directly supports choosing Cool over Archive when immediate access is needed.
Reference:
learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers#cool-access-tier and #archive-access-tier

Question: 14

HOTSPOT You have an on-premises file server that stores 2 TB of data files. You plan to move the data files to Azure Blob storage in the Central Europe region. You need to recommend a storage account type to store the data files and a replication solution for the storage account. The solution must meet the following requirements: Be available if a single Azure datacenter fails. Support storage tiers. Minimize cost. What should you recommend? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

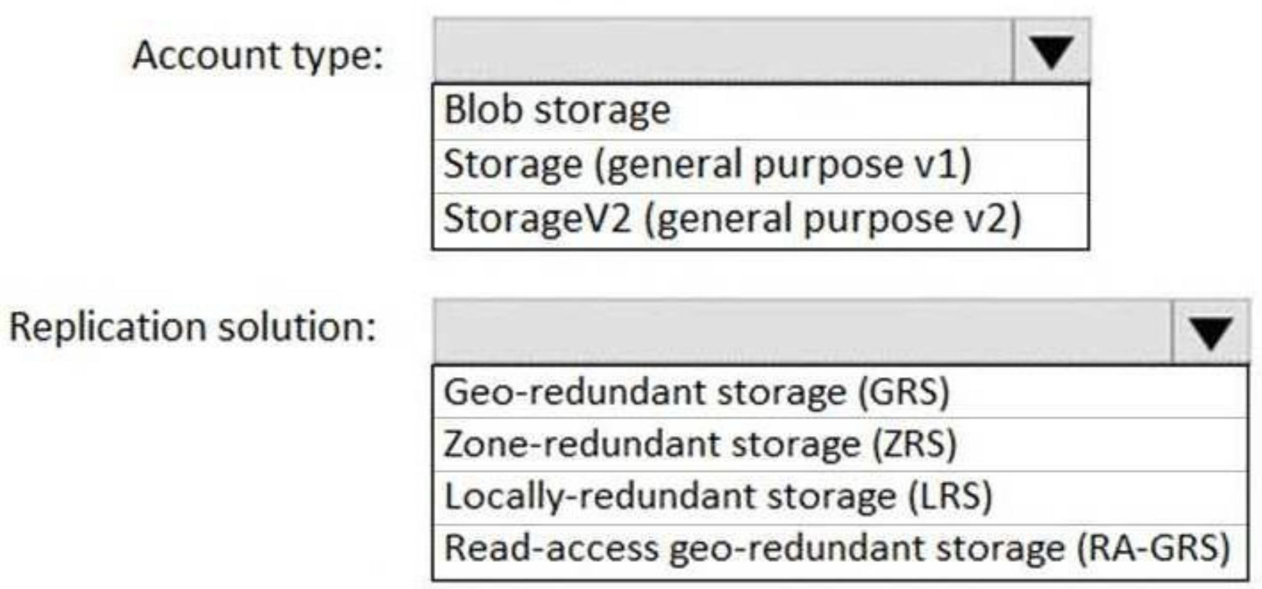
Account type:
Blob storage
Storage (general purpose v1)
StorageV2 (general purpose v2)
Replication solution:
Geo-redundant storage (GRS)
Zone-redundant storage (ZRS)
Locally-redundant storage (LRS)
Read-access geo-redundant storage (RA-GRS)

Answer:

Account type: StorageV2 (general purpose v2)

Replication solution: Zone-redundant storage (ZRS)

Explanation:

The solution requires an account type that supports storage tiers to help minimize costs, making StorageV2 (general purpose v2) the correct choice. The legacy Storage (general purpose v1) and Blob storage account types either lack support for all tiers or are not the recommended modern standard.

The solution must also be available if a single Azure datacenter fails.

- Locally-redundant storage (LRS) is insufficient as it only replicates data within a single datacenter.

- Zone-redundant storage (ZRS) meets the requirement by synchronously replicating data across three distinct datacenters (Availability Zones) within the primary region.

- Geo-redundant storage (GRS) and Read-access geo-redundant storage (RA-GRS) also meet the availability requirement but at a higher cost because they replicate data to a secondary region.

To minimize cost while meeting the availability requirement, ZRS is the most appropriate replication solution.

References:

Microsoft Documentation, Storage account overview, Azure Storage.
Section: "Types of storage accounts"
Details: This document explicitly states that General-purpose v2 accounts are recommended for most scenarios and support the latest Azure Storage features, including access tiers (Hot, Cool, Archive).
Microsoft Documentation, Data redundancy, Azure Storage.
Section: "Redundancy in the primary region" and "Redundancy in a secondary region"
Details: This documentation clarifies that Zone-redundant storage (ZRS) "protects your data against datacenter-level failures" by replicating it across three availability zones within the primary region. It also positions Geo-redundant storage (GRS) as a higher-cost option designed to protect against regional outages, which is beyond the stated requirement.

Question: 15

HOTSPOT Your company deploys several Linux and Windows virtual machines (VMs) to Azure. The VMs are deployed with the Microsoft Dependency Agent and the Log Analytics Agent installed by using Azure VM extensions. On-premises connectivity has been enabled by using Azure ExpressRoute. You need to design a solution to monitor the VMs. Which Azure monitoring services should you use? To answer, select the appropriate Azure monitoring services in the answer area. NOTE: Each correct selection is worth one point.

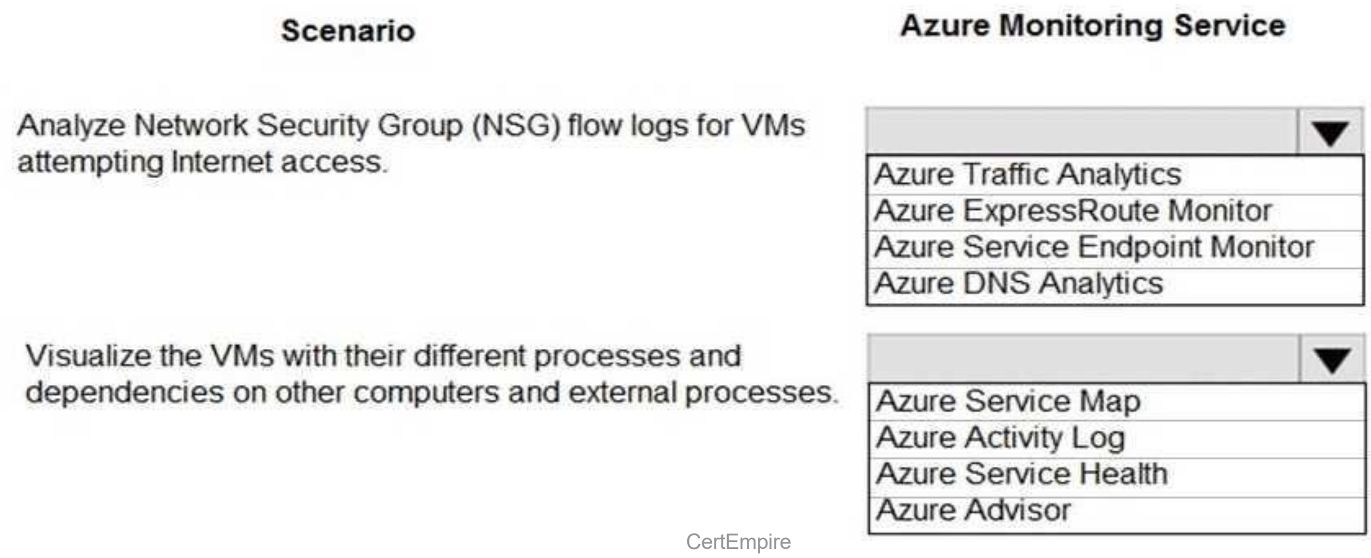

Scenario	Azure Monitoring Service
Analyze Network Security Group (NSG) flow logs for VMs attempting Internet access.	Azure Traffic Analytics Azure ExpressRoute Monitor Azure Service Endpoint Monitor Azure DNS Analytics
Visualize the VMs with their different processes and dependencies on other computers and external processes.	Azure Service Map Azure Activity Log Azure Service Health Azure Advisor

Answer:

Scenario 1: Azure Traffic Analytics

Scenario 2: Azure Service Map

Explanation:

Azure Traffic Analytics is the correct service for analyzing Network Security Group (NSG) flow logs. It processes these logs to provide insights into traffic flow, identify security threats, and understand network activity across Azure subscriptions. It aggregates traffic data and visualizes it, which directly addresses the need to analyze logs for VMs attempting internet access.

Azure Service Map is designed specifically to visualize VM processes and their dependencies. It automatically discovers application components on Windows and Linux systems and maps the communication between them. Its functionality relies on the Microsoft Dependency Agent, which the scenario explicitly states is installed on the VMs, making it the precise tool for visualizing these dependencies.

References:

Azure Traffic Analytics:

Microsoft Azure Documentation, "Traffic Analytics," Overview section. States, "Traffic analytics is a cloud-based solution that provides visibility into user and application activity in cloud networks. Traffic analytics analyzes Network Watcher network security group (NSG) flow logs to provide insights into traffic flow in your Azure cloud."

Azure Service Map:

Microsoft Azure Documentation, "Service Map," Overview section. Explains, "Service Map automatically discovers application components on Windows and Linux systems and maps the communication between services... It relies on the Dependency Agent to collect data about processes running on the virtual machine and their dependencies."

Question: 16

DRAG DROP You are designing a virtual machine that will run Microsoft SQL Server and will contain two data disks. The first data disk will store log files, and the second data disk will store dat a. Both disks are P40 managed disks. You need to recommend a caching policy for each disk. The policy must provide the best overall performance for the virtual machine. Which caching policy should you recommend for each disk? To answer, drag the appropriate policies to the correct disks. Each policy 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.

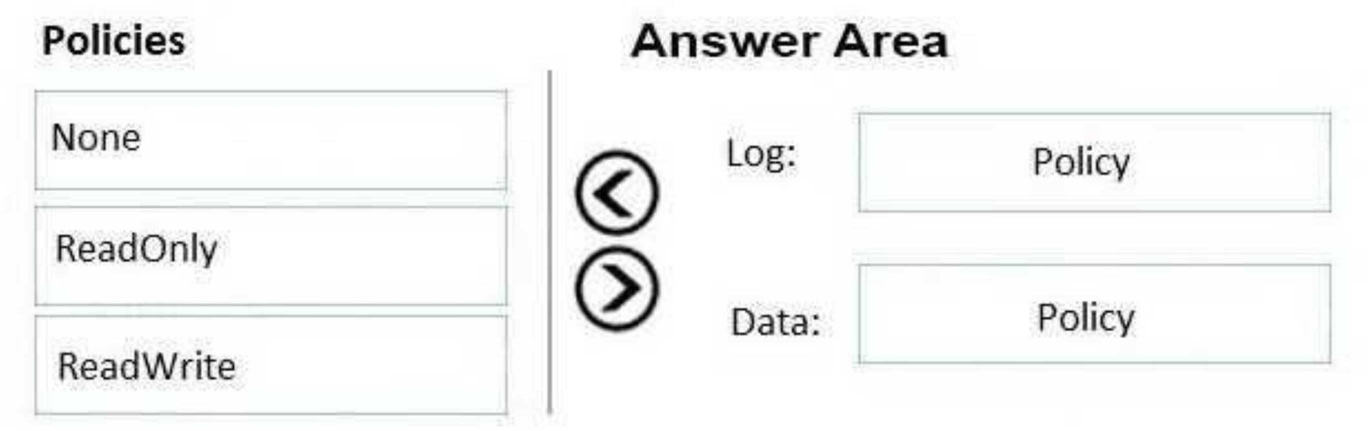

Policies

None

ReadOnly

ReadWrite

Answer Area

Log: Policy

Data: Policy

Answer:

Log: None

Data: ReadOnly

Explanation:

The optimal caching policy depends on the workload type for each disk.

- Log (Transaction Log Files): The SQL Server transaction log is write-intensive, and data integrity is paramount. Using any form of caching for the log file disk introduces the risk of data loss if the VM host were to crash before the cache is flushed to persistent storage. Therefore, disabling caching (None) is the best practice to ensure all log writes are committed directly and durably to the disk.

- Data (Database Files): SQL Server data files experience a high volume of read operations, especially for queries. Enabling ReadOnly caching significantly improves performance by serving frequently requested data from the faster host cache. This reduces the latency of read operations. Write caching is not recommended for data files as it can lead to data integrity issues.

References:

Microsoft Learn Performance best practices for SQL Server on Azure VMs: This official guide explicitly states the recommended caching policies.
Section/Area: Storage: Caching
Quote: "For data files, enable read caching. .... For log files, do not enable caching (None)." This document confirms that ReadOnly caching is recommended for data disks to accelerate reads, while caching should be disabled for the write-sensitive log disks to ensure data durability.
Microsoft Azure Docs Azure premium storage: design for high performance: This document details the workings of Azure disk caching.
Section/Area: Disk caching
Details: It explains that ReadOnly caching is optimal for read-heavy workloads, which is characteristic of SQL data files. It also advises against using ReadWrite caching for database applications like SQL Server that manage their own data durability, reinforcing the choice of None for transaction logs.

Question: 17

DRAG DROP You are planning an Azure solution that will host production databases for a high-performance application. The solution will include the following components: Two virtual machines that will run Microsoft SQL Server 2016, will be deployed to different data centers in the same Azure region, and will be part of an Always On availability group. SQL Server data that will be backed up by using the Automated Backup feature of the SQL Server IaaS Agent Extension (SQLIaaSExtension) You identify the storage priorities for various data types as shown in the following table.

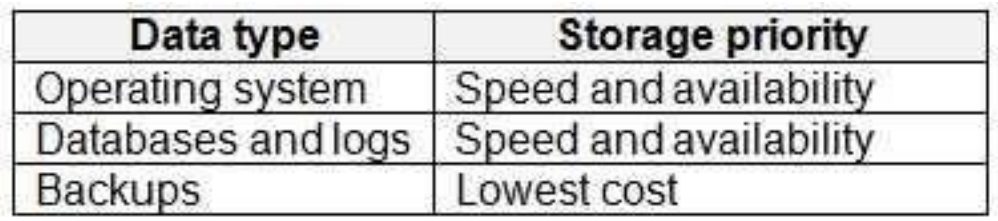

Data type	Storage priority
Operating system	Speed and availability
Databases and logs	Speed and availability
Backups	Lowest cost

Which storage type should you recommend for each data type? To answer, drag the appropriate storage types to the correct data types. Each storage 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.

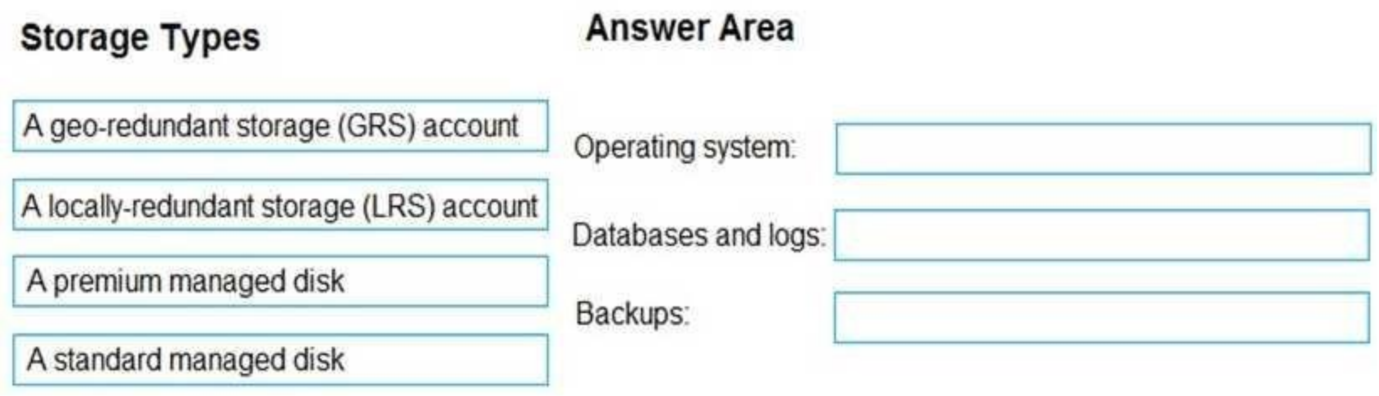

Storage Types

- A geo-redundant storage (GRS) account
- A locally-redundant storage (LRS) account
- A premium managed disk
- A standard managed disk

Answer Area

Operating system:

Databases and logs:

Backups:

Answer:

Operating system: A premium managed disk

Databases and logs: A premium managed disk

Backups: A locally-redundant storage (LRS) account

Explanation:

For both the Operating system and the Databases and logs, the key requirement is "Speed and availability". Premium managed disks are solid-state drives (SSDs) designed specifically for I/O-intensive production workloads like SQL Server, offering high throughput and low latency. This makes them the ideal choice to meet the performance demands.

For Backups, the primary driver is the "Lowest cost". The SQL Server IaaS Agent Extension stores backups in an Azure Storage account. Among the storage account redundancy options, Locally-redundant storage (LRS) is the least expensive. It maintains three copies of the data within a single data center, providing a durable and cost-effective solution suitable for backups where minimizing cost is the main priority.

References:

Azure Managed Disk Types: According to the official Microsoft documentation, Premium SSDs are recommended for "Production and performance sensitive SQL server" workloads and the Operating System Disk due to their high-performance characteristics.
Source: Microsoft Docs, "Select a disk type for Azure IaaS VMs - managed disks," Disk type comparison section.
Performance Best Practices for SQL Server on Azure VMs: The storage guidelines for SQL Server on Azure VMs explicitly recommend using Premium SSDs for both data and log files to ensure optimal performance.
Source: Microsoft Docs, "Performance best practices for SQL Server on Azure VMs," Storage section.
Azure Storage Redundancy Options: The documentation on data redundancy in Azure Storage clearly identifies Locally-redundant storage (LRS) as the "lowest-cost redundancy option" that offers the least durability compared to other options like GRS but still provides 11 nines of durability, making it suitable for cost-sensitive data like backups.
Source: Microsoft Docs, "Data redundancy - Azure Storage," Locally-redundant storage section.
SQL Server IaaS Agent Extension Backups: The documentation for the Automated Backup feature confirms that it uses an Azure storage account to store the backup files. The choice between LRS and GRS for this account depends on the trade-off between cost and disaster recovery requirements, and for the "lowest cost" priority, LRS is the correct selection.
Source: Microsoft Docs, "Automated Backup for SQL Server on Azure Virtual Machines," Prerequisites and Configuration sections.

Question: 18

HOTSPOT Your on-premises network contains a file server named Server1 that stores 500 GB of data. You need to use Azure Data Factory to copy the data from Server1 to Azure Storage. You add a new data factory. What should you do next? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

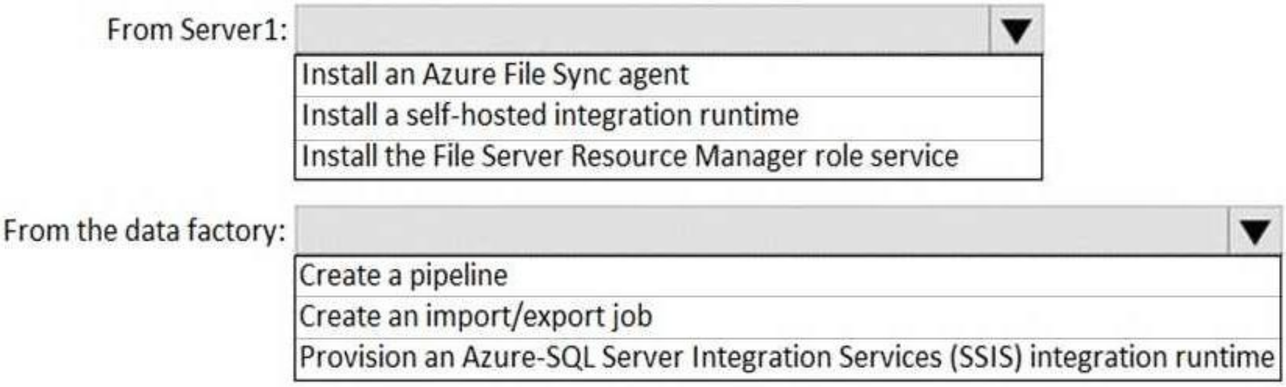

From Server1:

- Install an Azure File Sync agent
- Install a self-hosted integration runtime
- Install the File Server Resource Manager role service

From the data factory:

- Create a pipeline
- Create an import/export job
- Provision an Azure-SQL Server Integration Services (SSIS) integration runtime

Answer:

Correct Answer: Install a self-hosted integration runtime

Correct Answer: Create a pipeline

Explanation:

To enable Azure Data Factory, a cloud service, to access data sources located in an on-premises private network, a self-hosted integration runtime (SHIR) must be installed on a machine within that network. The SHIR acts as a secure gateway, executing data movement and transformation activities locally and dispatching them to Azure. Without the SHIR, the cloud-based Data Factory service has no way to connect to and pull data from Server1.

In Azure Data Factory, a pipeline is a logical grouping of activities that perform a task. The core task here is to copy data. This is accomplished by a "Copy Data" activity, which must be contained within a pipeline. After creating the Data Factory instance and setting up the integration runtime, the next step is to define this workflow by creating a pipeline and adding the necessary activities to it.

References:

Microsoft Documentation, Azure Data Factory and Synapse Analytics, "Create and configure a self-hosted integration runtime". This document states, "The self-hosted integration runtime is the compute infrastructure that Azure Data Factory and Synapse pipelines use to provide data integration capabilities across different network environments... For data movement, the self-hosted integration runtime moves data between a data source and a destination."

Microsoft Documentation, Azure Data Factory and Synapse Analytics, "Pipelines and activities in Azure Data Factory and Azure Synapse Analytics". This resource defines a pipeline as "a logical grouping of activities that together perform a task." It further clarifies that "An activity in a pipeline defines an action to perform on your data. For example, you can use a copy activity to copy data between data stores."
Microsoft Documentation, Azure Data Factory and Synapse Analytics, "Tutorial: Copy data from an on-premises data store to a cloud data store by using Azure Data Factory". This tutorial outlines the end-to-end process. The high-level steps after creating a data factory are: 1. Create a self-hosted integration runtime. 2. Create linked services. 3. Create datasets. 4. Create a pipeline. This confirms that installing the SHIR and creating the pipeline are the immediate subsequent steps.

Question: 19

You use Azure virtual machines to run a custom application that uses an Azure SQL database on the back end. The IT apartment at your company recently enabled forced tunneling, Since the configuration change, developers have noticed degraded performance when they access the database You need to recommend a solution to minimize latency when accessing the database. The solution must minimize costs What should you include in the recommendation?

A. Azure SQL Database Managed instance

B. Azure virtual machines that run Microsoft SQL Server servers

C. Always On availability groups

D. virtual network (VNET) service endpoint

Answer:

D

Explanation:

The performance degradation is caused by forced tunneling, which routes all internet-bound traffic from the Azure virtual machines back to the on-premises network before it reaches the public endpoint of the Azure SQL Database. This round-trip introduces significant latency.

A virtual network (VNet) service endpoint for Microsoft.Sql resolves this by providing a direct, optimized route from the VNet to the Azure SQL Database service over the Azure backbone network. This traffic bypasses the forced tunnel, eliminating the latency caused by hairpinning through the on-premises network. This solution directly addresses the performance issue and is the most cost-effective, as there is no additional charge for using VNet service endpoints.

Why Incorrect Options are Wrong:

A. Azure SQL Database Managed Instance: This is a more expensive service tier and a significant architectural change, which violates the "minimize costs" requirement.

B. Azure virtual machines that run Microsoft SQL Server servers: Migrating from a PaaS service to IaaS increases management overhead and is not a cost-effective solution for a network routing problem.

C. Always On availability groups: This is a high-availability and disaster recovery feature; it does not address network latency caused by forced tunneling.

References:

1. Microsoft Learn Virtual Network service endpoints: This document explains how service endpoints provide an optimized route for Azure service traffic. "Endpoints always take service traffic directly from your virtual network to the service on the Microsoft Azure backbone network." This direct path bypasses any user-defined routes (UDRs) or BGP routes that implement forced

tunneling.
Source: Microsoft Azure Documentation, "Virtual Network service endpoints overview", Section: "Key benefits".

2. Microsoft Learn Use virtual network service endpoints and rules for servers: This document details the implementation for Azure SQL Database. "By adding a virtual network rule, you're enabling a service endpoint for Azure SQL Database from the subnet... All traffic to Azure SQL Database from that subnet is then routed through the Azure private backbone."
Source: Microsoft Azure Documentation, "Use virtual network service endpoints and rules for servers in Azure SQL Database", Section: "Implementation steps".

3. Microsoft Learn Azure SQL Database and Azure Synapse Analytics connectivity architecture: This document describes how traffic flows. In the "Connectivity from within Azure" section, it explains that when connecting from a VNet with service endpoints, the connection is established over the Azure network backbone, avoiding the public internet gateway.
Source: Microsoft Azure Documentation, "Azure SQL Database and Azure Synapse Analytics connectivity architecture", Section: "Connectivity from within Azure".

Question: 20

Your network contains an on-premises Active Directory forest. You discover that when users change jobs within your company, the membership of the user groups are not being updated. As a result, the users can access resources that are no longer relevant to their job. You plan to integrate Active Directory and Azure Active Directory (Azure AD) by using Azure AD Connect. You need to recommend a solution to ensure that group owners are emailed monthly about the group memberships they manage. What should you include in the recommendation?

A. conditional access policies

B. Tenant Restrictions

C. Azure AD access reviews

D. Azure AD Identity Protection

Answer:

C

Explanation:

Azure AD access reviews are a feature of Azure AD Identity Governance designed specifically to address this scenario. They enable organizations to manage group memberships, access to enterprise applications, and privileged role assignments. You can configure recurring access reviews (e.g., monthly) that automatically prompt group owners via email to review and attest to their members' continued need for access. This process helps reduce the risk associated with excessive access permissions by ensuring that only the right people have continued access to resources.

Why Incorrect Options are Wrong:

A. conditional access policies: These policies enforce access controls based on signals like user location or device state, but they do not manage or review the lifecycle of group memberships.
B. Tenant Restrictions: This feature controls user access to external Azure AD tenants from your network or devices; it is not used for managing internal group memberships.
D. Azure AD Identity Protection: This service detects and remediates identity-based risks and vulnerabilities but does not include a feature for periodic attestation of group memberships by owners.

References:

1. Azure AD access reviews:
Microsoft Documentation, Azure Active Directory, "What are Azure AD access reviews?". States, "Azure Active Directory (Azure AD) access reviews enable organizations to efficiently manage group memberships, access to enterprise applications, and role assignments. User's access can

be reviewed on a regular basis to make sure only the right people have continued access."
Microsoft Documentation, Azure Active Directory, "Create an access review of groups and applications in Azure AD access reviews". Under the "Reviewers" section, it details how to select "Group owners" as the reviewers for the access review.

2. Conditional Access policies:

Microsoft Documentation, Azure Active Directory, "What is Conditional Access?". Explains, "Conditional Access policies at their simplest are if-then statements, if a user wants to access a resource, then they must complete an action."

3. Tenant Restrictions:

Microsoft Documentation, Azure Active Directory, "Use tenant restrictions to manage access to SaaS cloud applications". Describes the feature as a way to "control access to SaaS cloud applications, based on the Azure AD tenant the applications use for single sign-on."

4. Azure AD Identity Protection:

Microsoft Documentation, Azure Active Directory, "What is Identity Protection?". Defines the tool as one that "allows organizations to accomplish three key tasks: Automate the detection and remediation of identity-based risks, Investigate risks using data in the portal, Export risk detection data..."

Question: 21

HOTSPOT You plan to create an Azure environment that will contain a root management group and 10 child management groups. Each child management group will contain five Azure subscriptions. You plan to have between 10 and 30 resource groups in each subscription. You need to design an Azure governance solution. The solution must meet the following requirements: • Use Azure Blueprints to control governance across all the subscriptions and resource groups. • Ensure that Blueprints-based configurations are consistent across all the subscriptions and resource groups. • Minimize the number of blueprint definitions and assignments. What should you include in the solution? To answer, select the appropriate options in the answer are a. NOTE: Each correct selection is worth one point.

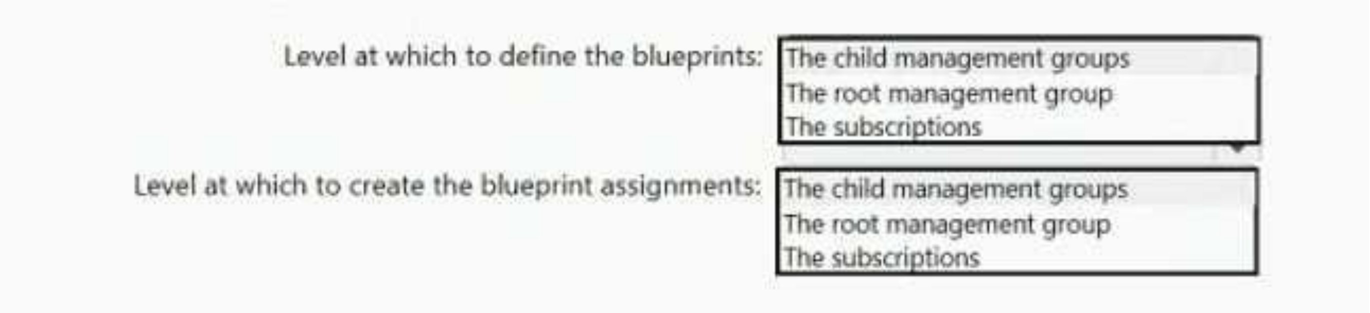

Level at which to define the blueprints:

The child management groups
The root management group
The subscriptions

Level at which to create the blueprint assignments:

The child management groups
The root management group
The subscriptions

Answer:

Level at which to define the blueprints: The root management group

Level at which to create the blueprint assignments: The root management group

Explanation:

To ensure governance is consistent across all subscriptions while minimizing administrative effort, both the blueprint definition and its assignment should be made at the highest possible level in the management group hierarchy.

• Blueprint Definition: Defining the blueprint at the root management group makes it available to all child management groups and subscriptions within the hierarchy. This creates a single, reusable source of truth, satisfying the need for consistency and minimizing the number of definitions to just one.

• Blueprint Assignment: Assigning the blueprint at the root management group applies it to all nested management groups and subscriptions. This single assignment ensures that all current and future subscriptions under the root inherit the same governance controls, meeting the requirement to minimize assignments.

This approach is the most efficient method for applying a consistent baseline configuration across the entire Azure environment as described.

References:

Microsoft Azure Documentation, "What is Azure Blueprints?":
Section: "Blueprint definition"
Content: "You can save your blueprint definition in a management group or subscription. If you save it to a management group, it can be assigned to any child subscription of that management group. This feature allows you to only have to define the blueprint once and reuse it for control and consistency at scale." This supports defining the blueprint at the highest level (the root management group) for maximum reuse and consistency.
Microsoft Azure Documentation, "Understand the lifecycle of an Azure Blueprint":
Section: "Blueprint assignment"
Content: "A blueprint can be assigned to a management group or subscription." It further explains that assignments are inherited by resources within the assigned scope. Assigning to the root management group is the most efficient way to ensure all subscriptions inherit the blueprint's configuration, thereby minimizing the number of individual assignments.

Question: 22

HOTSPOT You have five .NET Core applications that run on 10 Azure virtual machines in the same subscription. You need to recommend a solution to ensure that the applications can authenticate by using the same Azure Active Directory (Azure AD) identity. The solution must meet the following requirements: Ensure that the applications can authenticate only when running on the 10 virtual machines. Minimize administrative effort. 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.

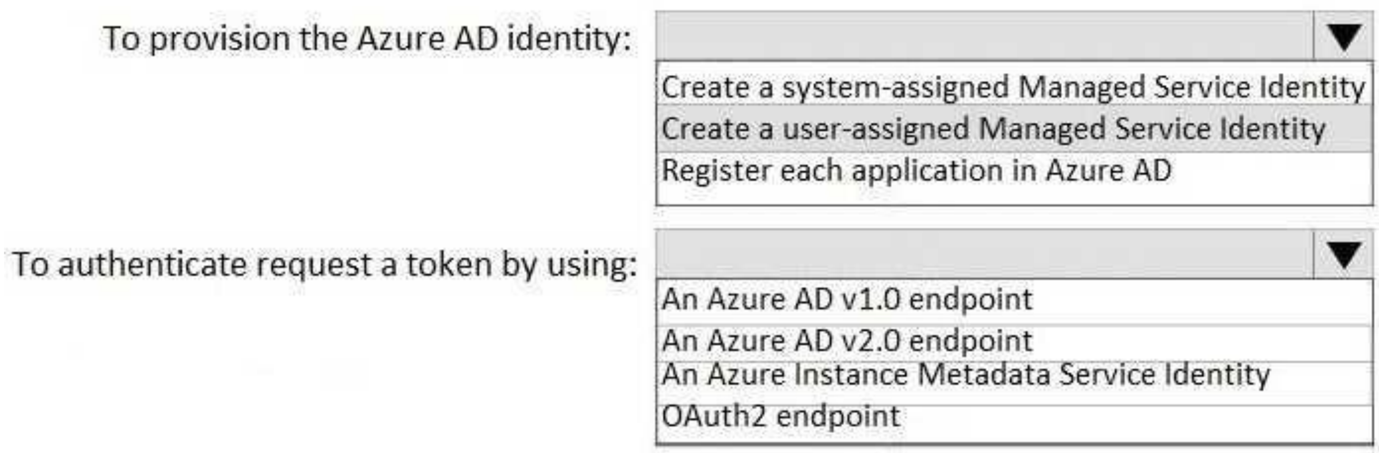

To provision the Azure AD identity:

- Create a system-assigned Managed Service Identity
- Create a user-assigned Managed Service Identity
- Register each application in Azure AD

To authenticate request a token by using:

- An Azure AD v1.0 endpoint
- An Azure AD v2.0 endpoint
- An Azure Instance Metadata Service Identity
- OAuth2 endpoint

Answer:

To provision the Azure AD identity: Create a user-assigned Managed Service Identity

To authenticate request a token by using: An Azure Instance Metadata Service Identity endpoint

Explanation:

A user-assigned managed identity is the correct choice because it is a standalone Azure resource that can be created once and then assigned to multiple Azure resources, in this case, the 10 virtual machines. This approach allows all applications on those VMs to share the same identity, satisfying the core requirements while minimizing administrative effort.

Applications running on a VM with a managed identity acquire authentication tokens by making a request to the Azure Instance Metadata Service (IMDS) identity endpoint. This is a local, non-routable endpoint (http://169.254.169.254/metadata/identity/oauth2/token) accessible only from within the VM. This mechanism ensures that only applications running on the specified 10 VMs can request tokens, fulfilling the security constraint.

References:

Microsoft. (n.d.). Managed identities for Azure resources. Microsoft Learn. Retrieved from learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview.
Reference Point: In the section "Managed identity types," the documentation states, "A user-assigned managed identity is created as a standalone Azure resource... It can be assigned to one or more Azure resources." This supports using a user-assigned identity to be shared across the 10 VMs.
Microsoft. (n.d.). How to use managed identities for Azure resources on an Azure VM to acquire an access token. Microsoft Learn. Retrieved from learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token.
Reference Point: Under the "Acquire a token using REST API" section, it specifies, "To get a token, your code makes a REST call to the Azure Instance Metadata Service (IMDS), which is accessible at a well-known, non-routable IP address (169.254.169.254)... The VM must be running and have networking correctly configured to make the request." This confirms the use of the IMDS endpoint for token acquisition.

Question: 23

HOTSPOT You plan to deploy a network-intensive application to several Azure virtual machines. You need to recommend a solution that meets the following requirements: Minimizes the use of the virtual machine processors to transfer data Minimizes network latency Which virtual machine size and feature should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

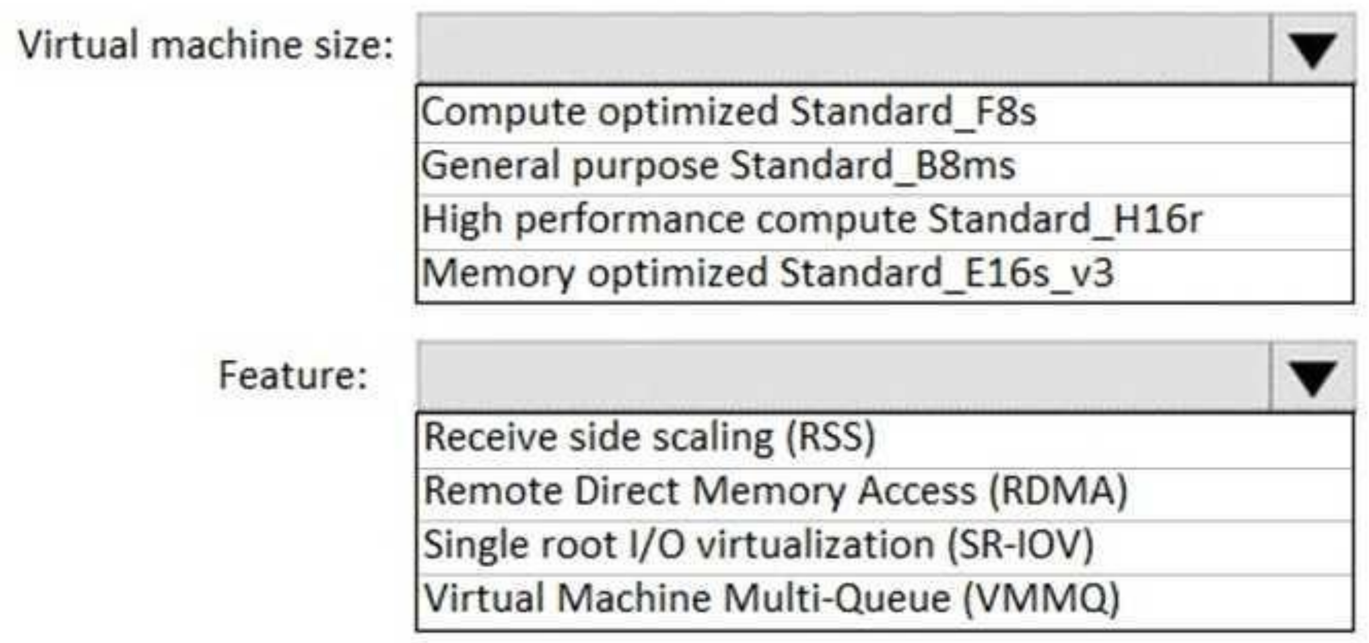

Virtual machine size: ▼

Compute optimized Standard_F8s
General purpose Standard_B8ms
High performance compute Standard_H16r
Memory optimized Standard_E16s_v3

Feature: ▼

Receive side scaling (RSS)
Remote Direct Memory Access (RDMA)
Single root I/O virtualization (SR-IOV)
Virtual Machine Multi-Queue (VMMQ)

Answer:

Virtual machine size: High performance compute StandardH16r

Feature: Remote Direct Memory Access (RDMA)

Explanation:

To minimize network latency and processor overhead for data transfers, the ideal solution involves a virtual machine and feature designed for High-Performance Computing (HPC).

The StandardH16r is an H-series Azure VM size specifically engineered for HPC workloads. The 'r' suffix indicates that it supports Remote Direct Memory Access (RDMA) over a high-speed InfiniBand network.

RDMA technology allows network adapters to transfer data directly to and from application memory on different virtual machines, bypassing the CPU, OS kernel, and TCP/IP stack. This direct memory access path drastically reduces latency and frees up CPU cycles, directly satisfying both of the stated requirements for the network-intensive application.

References:

Microsoft Azure Documentation, "H-series virtual machine sizes": This document explicitly states that H-series VMs are optimized for high-performance computing scenarios. It confirms that specific sizes, including the H16r, are RDMA-capable.
Section: "H-series" and "H-series feature comparison table".
Microsoft Azure Documentation, "Enable InfiniBand": This source details how RDMA-capable VMs, such as the H-series, communicate over the low-latency InfiniBand network.
Section: "Overview" and "RDMA capable VM instances". It states, "RDMA (Remote Direct Memory Access) provides low-latency, high-throughput networking for applications...".
Microsoft Azure Documentation, "Sizes for virtual machines in Azure": This general documentation provides an overview of VM families, categorizing the H-series for "High performance compute."
Section: "High performance compute".

Question: 24

You deploy two instances of an Azure web app. One instance is in the East US Azure region and the other instance is in the West US Azure region. The web app uses Azure Blob storage to deliver large files to end users. You need to recommend a solution for delivering the files to the users. The solution must meet the following requirements: Ensure that the users receive files from the same region as the web app that they access. Ensure that the files only need to be updated once. Minimize costs. What should you include in the recommendation?

A. Azure File Sync

B. Distributed File System (DFS)

C. read-access geo-redundant storage (RA-GRS)

D. geo-redundant storage (GRS)

Answer:

C

Explanation:

Read-access geo-redundant storage (RA-GRS) is the optimal solution. It replicates data from a primary region to a secondary region and, crucially, provides a read-only endpoint for that secondary location. This allows the web app in the secondary region (West US) to serve files from its local replica, minimizing latency and meeting the regional access requirement. Writes are directed only to the primary region and replicated automatically by Azure, satisfying the "update once" requirement. As a built-in storage redundancy option, RA-GRS is a cost-effective way to achieve these goals without additional services.

Why Incorrect Options are Wrong:

A. Azure File Sync is designed to centralize file shares in Azure Files and synchronize them with on-premises Windows Servers, which is not the scenario described.
B. Distributed File System (DFS) is a Windows Server role. Implementing it in Azure would require virtual machines, adding unnecessary cost and management overhead.
D. geo-redundant storage (GRS) replicates data to a secondary region for disaster recovery but does not allow read access to that data unless a failover occurs.

References:

1. Microsoft Docs, "Azure Storage redundancy." Under the section "Read-access geo-redundant storage (RA-GRS)," it states, "Read-access geo-redundant storage (RA-GRS) maximizes availability for your storage account... RA-GRS provides read-only access to the data in the secondary location, in addition to geo-replication across two regions." This directly supports the regional read requirement.

2. Microsoft Docs, "Design highly available applications using RA-GRS." The section "Reading data from the secondary endpoint" explains the pattern: "If you configure your storage account for read access to the secondary region, then you can design your application to read data from the secondary endpoint if the primary endpoint becomes unavailable for any reason... you can design your application to send read requests to the secondary endpoint." This confirms the solution's architecture.
3. Microsoft Docs, "Data redundancy." The summary table in this document clearly shows that for GRS, the secondary region is "Not available for read or write access unless there is a failover," while for RA-GRS, it is "Available for read-only access." This highlights why GRS is incorrect and RA-GRS is correct.

Question: 25

A company needs a datastore created in Azure for an application. Below are the key requirements for the data store. Ability to store JSON based items Ability to use SQL like queries on the datastore Ability to provide low latency access to data items Which of the following would you consider as the data store?

A. Azure BLOB storage

B. Azure CosmosDB

C. Azure HDInsight

D. Azure Redis

Answer:

B

Explanation:

Azure Cosmos DB is the ideal choice as it uniquely satisfies all three requirements. It is a multi-model NoSQL database service that natively stores data as schema-agnostic JSON items. The Core (SQL) API provides a familiar SQL-like query language to interact with the JSON data. Furthermore, Azure Cosmos DB is designed for global distribution and guarantees single-digit-millisecond read and write latencies at the 99th percentile, backed by service level agreements (SLAs), fulfilling the low-latency access requirement for modern applications.

Why Incorrect Options are Wrong:

A. Azure BLOB storage is an object store for unstructured data like files and media. It lacks a native SQL query engine for JSON content and is not optimized for low-latency database operations.
C. Azure HDInsight is a cloud analytics service for processing large datasets. It is designed for high-throughput batch workloads, not for low-latency transactional access to individual data items.
D. Azure Redis is an in-memory key-value store, primarily used as a cache. While it provides extremely low latency, it does not support a SQL-like query language.

References:

1. Microsoft Documentation, "Welcome to Azure Cosmos DB": "Azure Cosmos DB is a fully managed NoSQL...database for modern app development. It offers single-digit millisecond response times... The native data model of items within an Azure Cosmos DB container is the JSON model." This document confirms the JSON storage and low-latency capabilities.
2. Microsoft Documentation, "SQL queries in Azure Cosmos DB for NoSQL": "You can query data using the SQL... The items in Azure Cosmos DB are stored as JSON. The SQL query language works with JSON data models..." This source validates the use of SQL-like queries on JSON

items.
3. Microsoft Documentation, "Introduction to Azure Blob Storage": "Azure Blob Storage is Microsoft's object storage solution for the cloud... optimized for storing massive amounts of unstructured data." This reference clarifies that Blob storage is not a database designed for structured queries.
4. Microsoft Documentation, "What is Azure HDInsight?": "Azure HDInsight is a managed, full-spectrum, open-source analytics service in the cloud for enterprises." This positions HDInsight as an analytics platform, not a low-latency transactional datastore.
5. Microsoft Documentation, "What is Azure Cache for Redis?": "Azure Cache for Redis is based on the popular open-source in-memory datastore, Redis. It's typically used as a cache to improve the performance and scalability of systems..." This source confirms its primary use as a cache and its foundation on Redis, which does not use SQL.

Question: 26

You have to design a Data Engineering solution for your company. The company currently has an Azure subscription. They also have application data hosted in a database on a Microsoft SQL Server hosted in their on-premises data center server. They want to implement the following requirements Transfer transactional data from the on-premises SQL server onto a data warehouse in Azure. Data needs to be transferred every day in the night as a scheduled job A managed Spark cluster needs to be in place for data engineers to perform analysis on the data stored in the SQL data warehouse. Here the data engineers should have the ability to develop notebooks in Scale, R and Python. They also need to have a data lake store in place for the ingestion of data from multiple data sources Which of the following would the use for hosting the data warehouse in Azure?

A. Azure Data Factory

B. Azure Databricks

C. Azure Data Lake Gen2 Storage accounts

D. Azure Synapse Analytics

Answer:

D

Explanation:

The core requirement is to host a data warehouse in Azure. Azure Synapse Analytics is an integrated analytics service that provides enterprise data warehousing and Big Data analytics. Its dedicated SQL pool (formerly Azure SQL Data Warehouse) is a massively parallel processing (MPP) system specifically designed for running complex queries on large volumes of data, which is the primary function of a data warehouse. While the scenario also mentions requirements for data integration (pipelines), a Spark cluster, and a data lake, Azure Synapse Analytics is a unified platform that can fulfill all these needs, with the dedicated SQL pool being the specific data warehousing component.

Why Incorrect Options are Wrong:

A. Azure Data Factory is a cloud-based ETL and data integration service for orchestrating data movement and transformation; it is not a data warehouse itself.
B. Azure Databricks is a managed Apache Spark platform for big data analytics and machine learning, not a relational data warehouse for structured data querying.
C. Azure Data Lake Gen2 is a scalable storage solution for big data analytics but lacks the MPP query engine and management features of a dedicated data warehouse.

References:

1. Microsoft Learn. (2023). What is Azure Synapse Analytics? Microsoft Azure Documentation. Retrieved from https://learn.microsoft.com/en-us/azure/synapse-analytics/overview-what-is. Reference Section: "Overview". The document states, "Azure Synapse Analytics is an enterprise analytics service that accelerates time to insight across data warehouses and big data systems. Azure Synapse brings together the best of SQL technologies used in enterprise data warehousing..." This directly identifies Synapse as the primary service for data warehousing.

2. Microsoft Learn. (2023). What is Azure Data Factory? Microsoft Azure Documentation. Retrieved from https://learn.microsoft.com/en-us/azure/data-factory/introduction. Reference Section: "Overview". This source defines Data Factory as "the cloud-based ETL and data integration service," distinguishing it from a data storage and analytics platform like a data warehouse.

3. Microsoft Learn. (2023). Introduction to Azure Data Lake Storage Gen2. Microsoft Azure Documentation. Retrieved from https://learn.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-introduction. Reference Section: "Overview". This document describes Azure Data Lake Storage Gen2 as "a set of capabilities dedicated to big data analytics, built on Azure Blob Storage," positioning it as a storage foundation, not a complete data warehouse solution with a query engine.

Question: 27

Your company currently has an application that is hosted on their on-premises environment. The application currently connects to two databases in the on-premises environment. The databases are named whizlabdb1 and whizlabdb2. You have to move the databases onto Azure. The databases have to support server-side transactions across both of the databases. Solution: You decide to deploy the databases to an Azure SQL database-managed instance. Would this fulfill the requirement?

A. Yes

B. No

Answer:

A

Explanation:

The proposed solution is correct. Azure SQL Managed Instance is a platform-as-a-service (PaaS) offering designed for maximum compatibility with on-premises SQL Server. A key feature it supports is the ability to execute server-side transactions that span multiple databases hosted within the same managed instance. By deploying both whizlabdb1 and whizlabdb2 to a single Azure SQL Managed Instance, the application can continue to use cross-database transactions as it did on-premises, thus fulfilling the requirement.

Why Incorrect Options are Wrong:

B. This is incorrect because Azure SQL Managed Instance is specifically designed to support features common in on-premises SQL Server environments, including transactions that span multiple databases within the same instance.

References:

1. Microsoft Learn. (2023). Group transactions within a single instance using Azure SQL Managed Instance. In "Azure SQL Managed Instance documentation". Retrieved from https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/transact-sql-tsql-differences-sql-server#transactions-with-multiple-databases. The documentation states, "Unlike Azure SQL Database, SQL Managed Instance supports distributed transactions on databases within the instance."
2. Microsoft Learn. (2023). What is Azure SQL Managed Instance?. In "Azure SQL Managed Instance documentation". Retrieved from https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview. This overview highlights the high degree of compatibility with the on-premises SQL Server Database Engine, which includes native support for features like cross-database queries and transactions.

Question: 28

Your company has an on-premises Hyper-V cluster that contains 20 virtual machines. Some of the virtual machines are based on Windows and some in Linux. You have to migrate the virtual machines onto Azure. You have to recommend a solution that would be used to replicate the disks of the virtual machines to Azure. The solution needs to ensure that the virtual machines remain available when the migration of the disks is in progress. You decide to create an Azure storage account and then run AzCopy Would this fulfill the requirement?

A. Yes

B. No

Answer:

B

Explanation:

The proposed solution is incorrect. AzCopy is a command-line utility designed for copying files and blobs to or from Azure Storage accounts. It performs a point-in-time data transfer. It does not have the capability to perform continuous, block-level replication of a running virtual machine's disk. To migrate a live VM while it remains available, a solution that captures and replicates ongoing disk changes is required. The standard and recommended tool for this scenario is Azure Migrate, which uses a replication appliance to ensure data consistency and minimal downtime by continuously replicating changes from the on-premises VM to Azure until the final cutover.

Why Incorrect Options are Wrong:

A. Yes: This is incorrect because AzCopy cannot perform live replication of a running VM's disk. Using it would result in a stale and likely inconsistent disk image, failing to meet the requirement of migrating a live system without data loss.

References:

1. Microsoft Docs, Azure Migrate, "Migrate Hyper-V VMs to Azure with Azure Migrate": In the "Replicate Hyper-V VMs" section, the process is described: "After initial replication finishes, delta replication begins. Incremental changes to on-premises disks are periodically replicated to the replica disks in Azure." This highlights the continuous replication capability that AzCopy lacks.
2. Microsoft Docs, Azure Storage, "Get started with AzCopy": The introductory section states, "AzCopy is a command-line utility that you can use to copy blobs or files to or from a storage account." This defines its purpose as a data copy tool, not a live system replication service.
3. Microsoft Docs, Azure Migrate, "About Azure Migrate": The documentation states, "The Azure Migrate: Server Migration tool provides...Replication of on-premises machines to Azure." This explicitly identifies Azure Migrate as the service for VM replication.

Question: 29

You have an Azure virtual machine named VM1 that runs Windows Server 2019 and contains 500 GB of data files. You are designing a solution that will use Azure Data Factory to transform the data files, and then load the files to Azure Data Lake Storage What should you deploy on VM1 to support the design?

A. the self-hosted integration runtime

B. the Azure Pipelines agent

C. the On-premises data gateway

D. the Azure File Sync agent

Answer:

A

Explanation:

To allow Azure Data Factory (ADF) to access and process data files located on a virtual machine within a private network (such as an Azure Virtual Network), a self-hosted integration runtime (SHIR) must be installed on that machine or another machine within the same network. The SHIR acts as a secure compute bridge, enabling ADF to execute data movement and transformation activities against data sources that are not publicly accessible. In this scenario, the SHIR on VM1 will allow ADF to connect to the local file system, read the data files, and push them to the transformation pipeline.

Why Incorrect Options are Wrong:

B. the Azure Pipelines agent: This agent is used by Azure DevOps to run build and release jobs for CI/CD pipelines, not for data integration with Azure Data Factory.
C. the On-premises data gateway: This gateway is used by services like Power BI, Power Apps, and Logic Apps to connect to on-premises data, but Azure Data Factory uses its own specific component, the Integration Runtime.
D. the Azure File Sync agent: This agent synchronizes files between a Windows Server and an Azure Files share. It is not used to facilitate data movement or transformation by Azure Data Factory.

References:

1. Microsoft Documentation - Azure Data Factory and Synapse Analytics, "Integration runtime". This document explicitly states, "A self-hosted integration runtime can be used for... Running a Copy activity between a cloud data store and a data store in a private network." It also clarifies its use for dispatching transform activities against compute resources in an Azure virtual network.

Source: Microsoft Learn, "Integration runtime in Azure Data Factory and Azure Synapse Analytics", Section: "Self-hosted integration runtime".

2. Microsoft Documentation - Azure Data Factory and Synapse Analytics, "Create and configure a self-hosted integration runtime". This guide details the process and purpose of the SHIR. The introduction states, "The self-hosted integration runtime is the compute infrastructure that Azure Data Factory and Synapse pipelines use to provide data integration capabilities across different network environments."

Source: Microsoft Learn, "Create and configure a self-hosted integration runtime", Section: "Introduction".

3. Microsoft Documentation - Power BI, "What is an on-premises data gateway?". This document lists the services that use the on-premises data gateway, such as Power BI, Power Apps, and Logic Apps. Azure Data Factory is notably absent, reinforcing that it uses a different mechanism (the SHIR).

Source: Microsoft Learn, "What is an on-premises data gateway?", Section: "Gateway architecture".

Question: 30

You plan to deploy multiple instances of an Azure web app across several Azure regions. You need to design an access solution for the app. The solution must meet the following replication requirements; • Support rate limiting. • Balance requests between all instances. • Ensure that users can access the app in the event of a regional outage. Solution: You use Azure Traffic Manager to provide access to the app. Does this meet the goal?

A. Yes

B. No

Answer:

B

Explanation:

Azure Traffic Manager is a DNS-based traffic load balancer. It successfully meets the requirements to balance requests across regional instances and provide high availability through endpoint monitoring and failover. However, Traffic Manager operates at the DNS layer and does not inspect the actual traffic flowing to the application. Because it does not process the HTTP/S requests, it cannot natively implement rate limiting, which is a Layer 7 (application layer) function. To meet all requirements, a service like Azure Front Door, which provides global load balancing and includes a Web Application Firewall (WAF) for rate limiting, would be necessary.

Why Incorrect Options are Wrong:

A. Yes: This is incorrect because the solution fails to meet the mandatory requirement of supporting rate limiting, as Azure Traffic Manager is a DNS-level service.

References:

1. Microsoft Learn, "What is Traffic Manager?"
Section: Introduction
Content: "Azure Traffic Manager is a DNS-based traffic load balancer. This service allows you to distribute traffic to your public facing applications across the global Azure regions." This establishes that Traffic Manager works at the DNS level, resolving a domain name to an IP address, but does not process the subsequent traffic.
2. Microsoft Learn, "Load-balancing options"
Section: "Choose a load-balancing service" - "Feature comparison" table.
Content: The feature comparison table explicitly shows that Azure Traffic Manager does not support "Web Application Firewall (WAF)" capabilities, which include rate limiting. In contrast, services like Azure Front Door and Azure Application Gateway are shown to support WAF.
3. Microsoft Learn, "Rate limiting with Azure Front Door"

Section: "What is a web application firewall rate limit rule on Azure Front Door?"
Content: "A rate limit rule for Azure Front Door controls the number of requests allowed from a particular client IP address to your application during a one-minute duration." This document demonstrates that rate limiting is a feature of Azure Front Door's WAF, not Traffic Manager.

Question: 31

You plan to deploy multiple instances of an Azure web app across several Azure regions. You need to design an access solution for the app. The solution must meet the following replication requirements: • Support rate limiting • Balance requests between all instances. • Ensure that users can access the app in the event of a regional outage Solution: You use Azure Load Balancer to provide access to the app. Does this meet the goal?

A. Yes

B. No

Answer:

B

Explanation:

The proposed solution is incorrect. Azure Load Balancer is a regional, Layer 4 (TCP/UDP) service. It cannot meet the specified requirements. Firstly, it does not natively support rate limiting, which is a Layer 7 (HTTP/S) feature typically handled by services like Azure Application Gateway WAF or Azure Front Door. Secondly, as a regional service, a standard Azure Load Balancer cannot balance traffic across multiple Azure regions or provide automatic failover in the event of a regional outage. A global load balancing solution, such as Azure Front Door or Azure Traffic Manager, is required to route traffic across regions and ensure high availability during a regional failure.

Why Incorrect Options are Wrong:

A. Yes: This is incorrect because Azure Load Balancer is a regional Layer 4 service and lacks the required global routing, regional failover, and native rate-limiting capabilities.

References:

1. Azure Architecture Center - Load-balancing options. This document explicitly states, "For global routing, we recommend Azure Front Door." It also categorizes Azure Load Balancer as a Regional load balancer, contrasting it with Global options like Front Door and Traffic Manager, which are necessary for regional outage scenarios.
Source: Microsoft Learn, Azure Architecture Center. (2023). Load-balancing options. Section: "Azure load-balancing services".
2. Azure Load Balancer overview. This documentation confirms that "Azure Load Balancer operates at layer 4 of the Open Systems Interconnection (OSI) model" and is a regional resource, which means it cannot route traffic between regions.
Source: Microsoft Learn. (2023). What is Azure Load Balancer?. Section: "Introduction".
3. Web Application Firewall (WAF) rate limiting. This document details how rate limiting is a

feature of Azure Application Gateway WAF and Azure Front Door, not Azure Load Balancer. It states, "Rate limiting allows you to detect and block abnormally high levels of traffic from any client IP address."
Source: Microsoft Learn. (2023). Rate limiting on Azure Application Gateway. Section: "Overview".

Question: 32

You are developing a sales application that will contain several Azure cloud services and handle different components of a transaction. Different cloud services will process customer orders, billing, payment inventory, and shipping. You need to recommend a solution to enable the cloud services to asynchronously communicate transaction information by using XML messages. What should you include in the recommendation?

A. Azure Data Lake

B. Azure Notification Hubs

C. Azure Queue Storage

D. Azure Service Fabric

Answer:

C

Explanation:

Azure Queue Storage is a service for storing large numbers of messages that can be accessed from anywhere in the world. It is designed for building scalable, decoupled applications. In this scenario, the different cloud services (orders, billing, inventory) can communicate asynchronously by placing XML messages into a queue. The sending service adds a message and can continue its work, while the receiving service can retrieve and process the message when it is ready. This pattern effectively decouples the components, improving the application's overall reliability and scalability, which is ideal for handling different stages of a transaction.

Why Incorrect Options are Wrong:

A. Azure Data Lake is a scalable data storage and analytics service. It is designed for big data workloads, not for real-time, transactional messaging between services.

B. Azure Notification Hubs is a massively scalable mobile push notification engine. Its purpose is to send notifications to client applications on various platforms, not for backend service-to-service communication.

D. Azure Service Fabric is a distributed systems platform for building and deploying microservices. While you could build a messaging system on it, it is not the messaging service itself.

References:

1. Microsoft Documentation, "What is Azure Queue Storage?": "Azure Queue Storage is a service for storing large numbers of messages. You access messages from anywhere in the world via authenticated calls using HTTP or HTTPS. A queue message can be up to 64 KB in size. A queue may contain millions of messages, up to the total capacity limit of a storage account.

Queues are commonly used to create a backlog of work to process asynchronously."
Source: Microsoft Docs, Azure Storage Documentation, Queues.
2. Microsoft Documentation, "Storage queues and Service Bus queues - compared and contrasted": "Azure Queue Storage... provides a simple REST-based Get/Put/Peek interface, providing reliable, persistent messaging within and between services... Use Queue storage when you need to store over 80 gigabytes of messages in a queue and you want a simple, easy to use queue." This document highlights its use for decoupling application components for increased scalability and reliability.
Source: Microsoft Docs, Azure Architecture Center, Application integration.
3. Microsoft Documentation, "What is Azure Notification Hubs?": "Azure Notification Hubs provide an easy-to-use and scaled-out push engine that allows you to send notifications to any platform (iOS, Android, Windows, etc.) from any back-end (cloud or on-premises)."
Source: Microsoft Docs, Azure Notification Hubs Documentation, Overview.
4. Microsoft Documentation, "Introduction to Azure Data Lake Storage Gen2": "Azure Data Lake Storage Gen2 is a set of capabilities dedicated to big data analytics, built on Azure Blob Storage."
Source: Microsoft Docs, Azure Storage Documentation, Data Lake Storage.

Question: 33

Your company has the divisions shown in the following table.

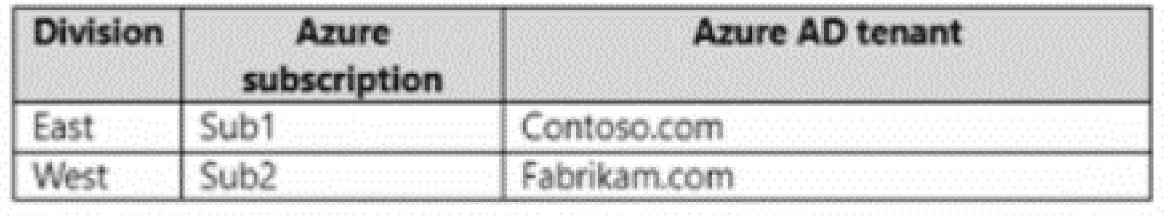

Division	Azure subscription	Azure AD tenant
East	Sub1	Contoso.com
West	Sub2	Fabrikam.com

Sub1 contains an Azure App Service web app named App1. Appl uses Azure AD for single-tenant user authentication. Users from contoso.com can authenticate to App1. You need to recommend a solution to enable users in the fabrikam.com tenant to authenticate to App1. What should you recommend?

A. Configure the Azure AD provisioning service.

B. Configure Supported account types in the application registration and update the sign-in endpoint.

C. Configure assignments for the fabrikam.com users by using Azure AD Privileged Identity Management (PIM).

D. Enable Azure AD pass-through authentication and update the sign-in endpoint

Answer:

B

Explanation:

The application is currently configured as a single-tenant app, which restricts authentication to users within its home tenant (contoso.com). To allow users from an external Azure AD tenant (fabrikam.com) to authenticate, the application must be reconfigured to be multi-tenant. This is accomplished by modifying the "Supported account types" setting in the application's registration within the Azure portal. Changing this setting to "Accounts in any organizational directory (Any Azure AD directory - Multitenant)" makes the application available to users from any Azure AD tenant. The application's sign-in endpoint logic must also be updated to handle requests from the generic /organizations or /common endpoint instead of the tenant-specific one.

Why Incorrect Options are Wrong:

A. The Azure AD provisioning service automates creating and managing user identities in other applications; it does not configure an application's authentication audience.
C. Azure AD Privileged Identity Management (PIM) is used to manage, control, and monitor access to privileged roles, not to enable standard cross-tenant user authentication.
D. Azure AD pass-through authentication is a sign-in method for hybrid identity that validates user

passwords against an on-premises Active Directory; it is not relevant for cross-tenant authentication.

References:

1. Microsoft Documentation: How to: Sign in any Azure Active Directory user using the multi-tenant application pattern.
Reference: In the section "Update the registration to be multi-tenant," the document states: "If you have an existing application and you want to make it multi-tenant, you need to open the application registration in the Azure portal and update Supported account types to Accounts in any organizational directory." This directly supports the chosen answer.
2. Microsoft Documentation: Quickstart: Register an application with the Microsoft identity platform.
Reference: In the "Register an application" section, step 4, "Supported account types," explicitly defines the option "Accounts in any organizational directory (Any Azure AD directory - Multitenant)" as the method to allow users with a work or school account from any organization to sign into the application.
3. Microsoft Documentation: Tenancy in Azure Active Directory.
Reference: The "App-level considerations" section explains the difference between single-tenant and multi-tenant applications. It clarifies that a multi-tenant application is "available to users in both its home tenant and other tenants." This conceptual document underpins the need to change the application's tenancy model to meet the requirement.

Question: 34

You need to design a highly available Azure SQL database that meets the following requirements: * Failover between replicas of the database must occur without any data loss. * The database must remain available in the event of a zone outage. * Costs must be minimized. Which deployment option should you use?

A. Azure SQL Database Premium

B. Azure SQL Database Hyperscale

C. Azure SQL Database Basic

D. Azure SQL Managed Instance Business Critical

Answer:

A

Explanation:

The Azure SQL Database Premium tier is the most appropriate choice. It supports zone-redundant configurations, which provision replicas in different availability zones within the same region. This architecture uses synchronous replication, ensuring that failovers occur with zero data loss (Recovery Point Objective - RPO=0) and that the database remains available during a zone-level outage. Compared to Hyperscale and Managed Instance Business Critical, the Premium tier provides these high-availability features at a lower cost, thus satisfying the "costs must be minimized" requirement for workloads that do not require the massive scale of Hyperscale or the instance-level features of Managed Instance.

Why Incorrect Options are Wrong:

B. Azure SQL Database Hyperscale: While it supports zone redundancy, this tier is designed for very large databases (VLDBs) and is not the most cost-effective option for general high-availability scenarios.
C. Azure SQL Database Basic: This tier does not support zone-redundant configurations and cannot meet the requirement to remain available during a zone outage.
D. Azure SQL Managed Instance Business Critical: This option meets the availability and data-loss requirements but is generally more expensive than Azure SQL Database Premium, failing the cost-minimization constraint.

References:

1. Microsoft Documentation, "High availability for Azure SQL Database and SQL Managed Instance": Under the "Zone-redundant availability" section, it states, "Zone-redundant configuration is available for databases in the... Premium, Business Critical, and Hyperscale service tiers... When you provision a database or an elastic pool with zone redundancy, Azure

SQL creates multiple synchronous secondary replicas in other availability zones." This confirms that Premium meets the zone outage and no data loss requirements.
2. Microsoft Documentation, "vCore purchasing model - Azure SQL Database": The "Premium service tier" section describes it as being designed for "I/O-intensive workloads that require high availability and low-latency I/O." The documentation confirms that zone redundancy is a configurable option for this tier.
3. Microsoft Documentation, "Service Tiers in the DTU-based purchase model": This document shows that the Basic tier has a "Basic availability" model with a single database file and is not designed for high availability or zone redundancy.
4. Microsoft Documentation, "Compare the vCore and DTU-based purchasing models of Azure SQL Database": This page highlights that the Premium tier (in both models) is designed for high performance and high availability, whereas Managed Instance is for "lift-and-shift of the largest number of SQL Server applications to the cloud with minimal changes," which often comes at a higher price point.

Question: 35

DRAG DROP You have an on-premises named App 1. Customers App1 to manage digital images. You plan to migrate App1 to Azure. You need to recommend a data storage solution for Appl. The solution must meet the following image storage requirements: Encrypt images at rest. Allow files up to 50M

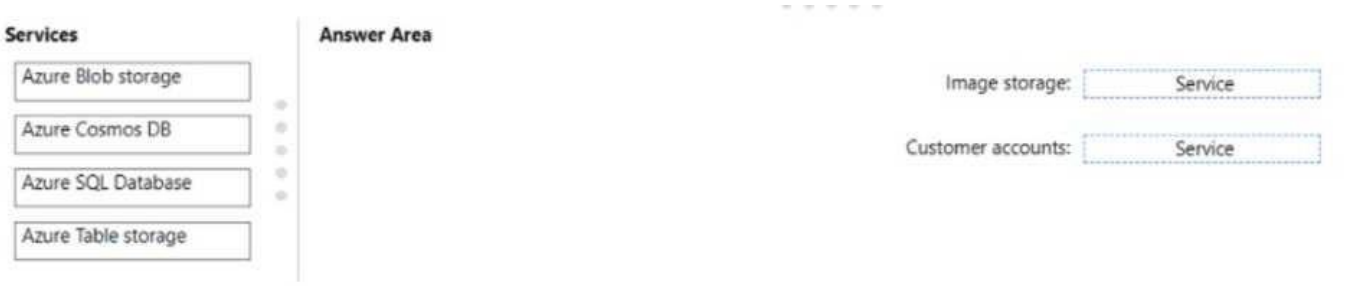

Services

Azure Blob storage

Azure Cosmos DB

Azure SQL Database

Azure Table storage

Answer Area

Image storage: Service

Customer accounts: Service

Answer:

Image storage: Azure Blob storage

Customer accounts: Azure SQL Database

Explanation:

Azure Blob storage is the optimal choice for image storage. It's specifically designed to store massive amounts of unstructured data, such as images, videos, and documents. It easily accommodates files up to 50 MB and provides server-side encryption by default, satisfying both requirements. Storing large binary files directly in a database is generally inefficient and not recommended.

Azure SQL Database is the most suitable service for customer accounts. Customer account data is typically structured and relational (e.g., user ID, name, email, password). As a fully managed relational database-as-a-service, Azure SQL Database provides transactional consistency, data integrity, and robust querying capabilities, which are essential for managing user account information effectively.

References:

Azure Blob Storage Documentation: Microsoft's official documentation states that Azure Blob storage is optimized for storing massive amounts of unstructured data. Common use cases include "Serving images or documents directly to a browser" and "Storing files for distributed access."

Source: Microsoft Docs, "Introduction to Azure Blob storage," Use cases section.

Azure SQL Database Documentation: The official documentation describes Azure SQL Database as a fully managed relational database service built for the cloud. It is ideal for applications that

require a relational data model with transactional consistency and data integrity, making it a standard choice for storing structured data like user profiles and customer accounts.
Source: Microsoft Docs, "What is Azure SQL Database?," Overview section.
Comparison of Azure Storage Options: Microsoft's "Choose a data storage approach in Azure" guide recommends Blob storage for "images, videos, documents...large binary objects" and relational databases like Azure SQL Database for "transactional data" and data requiring a "high degree of integrity," such as customer information.
Source: Microsoft Azure Architecture Center, "Choose a data storage approach in Azure," Relational databases and Blob storage sections.