AZ-305 Q&As

Designing Microsoft Azure Infrastructure Solutions

Pass Microsoft AZ-305 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/az-305.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

You are designing an app that will use Azure Cosmos DB to collate sales data from multiple countries. You need to recommend an API for the app. The solution must meet the following requirements:

1.

Support SQL queries.

2.

Support geo-replication.

3.

Store and access data relationally. Which API should you recommend?

A. PostgreSQL

B. NoSQL

C. Apache Cassandra

D. MongoDB

Correct Answer: A

QUESTION 2

You plan to migrate data to Azure.

The IT department at your company identifies the following requirements:

1.

The storage must support 1 PB of data.

2.

The data must be stored in blob storage.

3.

The storage must support three levels of subfolders.

4.

The storage must support access control lists (ACLs).

You need to meet the requirements.

What should you use?

A. a premium storage account that is configured for block blobs

B. a general purpose v2 storage account that has hierarchical namespace enabled

C. a premium storage account that is configured for page blobs

D. a premium storage account that is configured for files shares and supports large file shares

Correct Answer: B

Microsoft recommends that you use a GPv2 storage account for most scenarios. It supports up to 5 PB, and blob storage including Data Lake storage.

Note: A key mechanism that allows Azure Data Lake Storage Gen2 to provide file system performance at object storage scale and prices is the addition of a hierarchical namespace. This allows the collection of objects/files within an account to be organized into a hierarchy of directories and nested subdirectories in the same way that the file system on your computer is organized. With a hierarchical namespace enabled, a storage account becomes capable of providing the scalability and cost-effectiveness of object storage, with file system semantics that are familiar to analytics engines and frameworks.

References: https://docs.microsoft.com/en-us/azure/storage/common/storage-account-overview
https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-namespace

QUESTION 3

HOTSPOT

You are designing a data storage solution to support reporting.

The solution will ingest high volumes of data in the JSON format by using Azure Event Hubs. As the data arrives, Event Hubs will write the data to storage. The solution must meet the following requirements:

1.

Organize data in directories by date and time.

2.

Allow stored data to be queried directly, transformed into summarized tables, and then stored in a data warehouse.

3.

Ensure that the data warehouse can store 50 TB of relational data and support between 200 and 300 concurrent read operations. Which service should you recommend for each type of data store? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Data store for the ingested data:

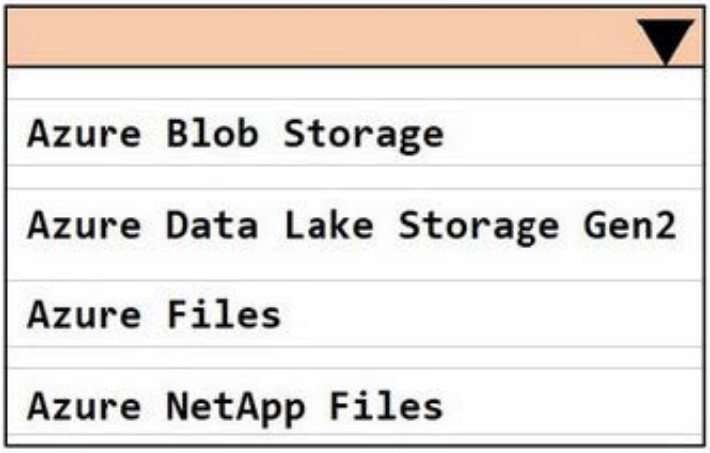

Data store for the data warehouse:

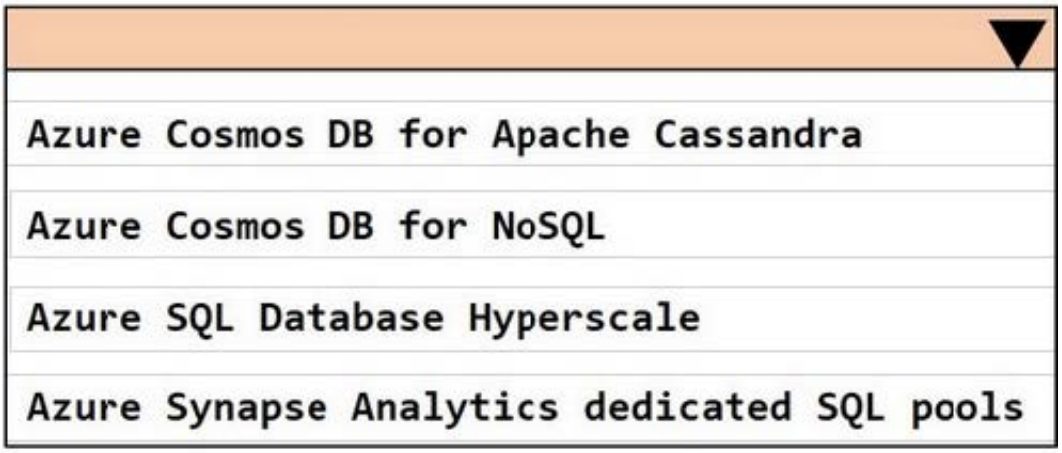

Correct Answer:

Data store for the ingested data:

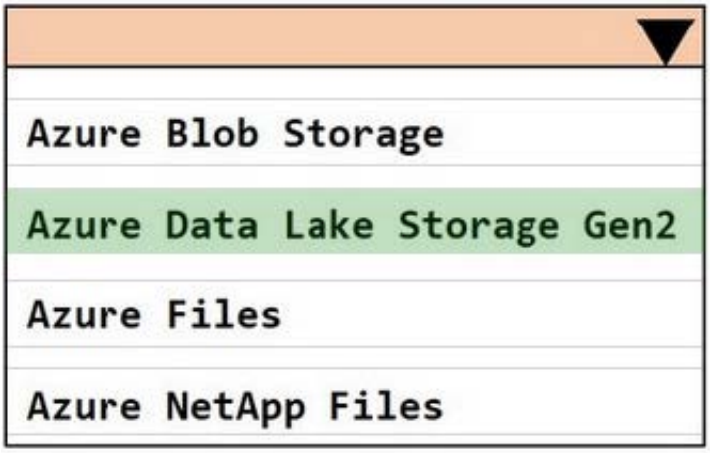

Data store for the data warehouse:

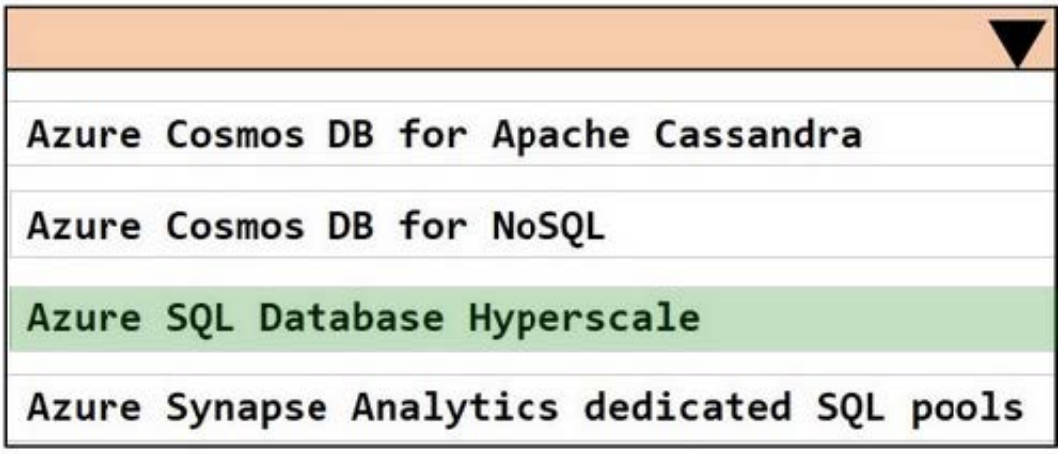

QUESTION 4

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 data.

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.

Select and Place:

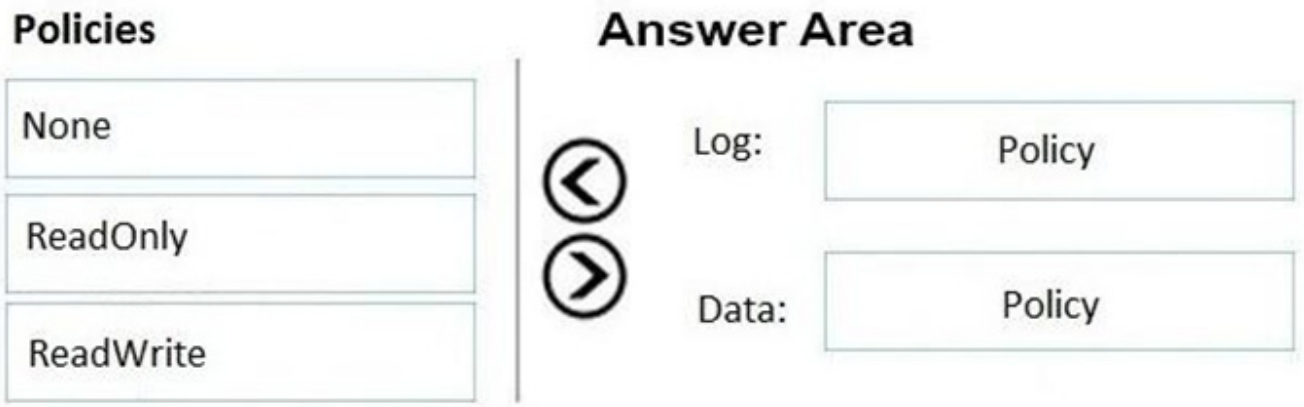

Correct Answer:

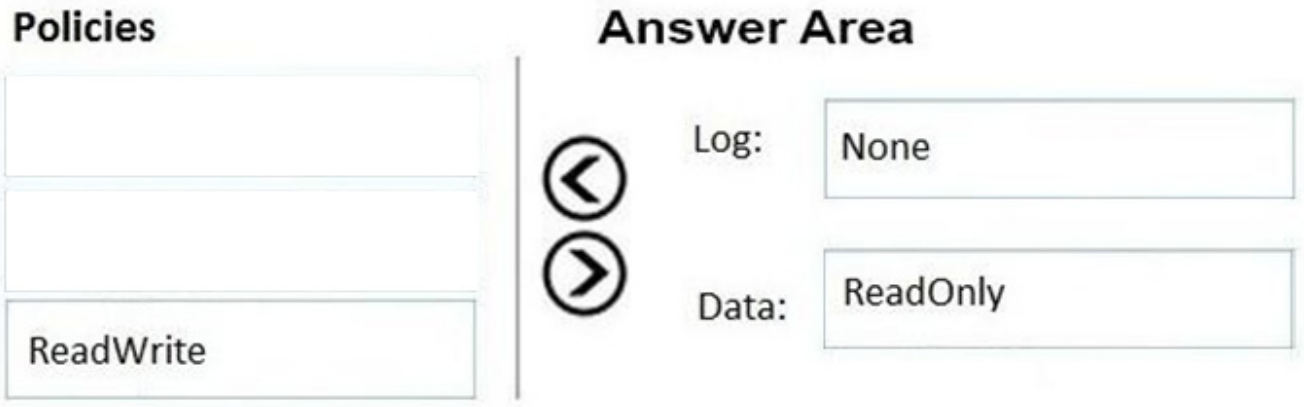

References: https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-sql-performance

QUESTION 5

A company named Contoso, Ltd. has an Azure Active Directory (Azure AD) tenant that is integrated with Microsoft Office 365 and an Azure subscription.

Contoso has an on-premises identity infrastructure. The infrastructure includes servers that run Active Directory Domain Services (AD DS), and Azure AD Connect

Contoso has a partnership with a company named Fabrikam, Inc. Fabrikam has an Active Directory forest and an Office 365 tenant. Fabrikam has the same on-premises identity infrastructure as Contoso.

A team of 10 developers from Fabrikam will work on an Azure solution that will be hosted in the Azure subscription of Contoso. The developers must be added to the Contributor role for a resource in the Contoso subscription.

You need to recommend a solution to ensure that Contoso can assign the role to the 10 Fabrikam developers. The solution must ensure that the Fabrikam developers use their existing credentials to access resources.

What should you recommend?

A. Configure a forest trust between the on-premises Active Directory forests of Contoso and Fabrikam.

B. Configure an organization relationship between the Office 365 tenants of Fabrikam and Contoso.

C. In the Azure AD tenant of Contoso, use MIM to create guest accounts for the Fabrikam developers.

D. Configure an AD FS relying party trust between the fabrikam and Contoso AD FS infrastructures.

Correct Answer: A

Trust configurations - Configure trust from managed forests(s) or domain(s) to the administrative forest. A one-way trust is required from production environment to the admin forest. Selective authentication should be used to restrict accounts in the admin forest to only logging on to the appropriate production hosts.

References:

https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/securing-privileged- access-reference-material

QUESTION 6

HOTSPOT

Your company has the divisions shown in the following table.

Division	Azure subscription	Azure Active Directory (Azure AD) tenant
East	Sub1, Sub2	East.contoso.com
West	Sub3, Sub4	West.contoso.com

You plan to deploy a custom application to each subscription. The application will contain the following:

1.

A resource group

2.

An Azure web app

3.

Custom role assignments

4.

An Azure Cosmos DB account

You need to use Azure Blueprints to deploy the application to each subscription.

What is the minimum number of objects required to deploy the application? To answer, select the appropriate options in

the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

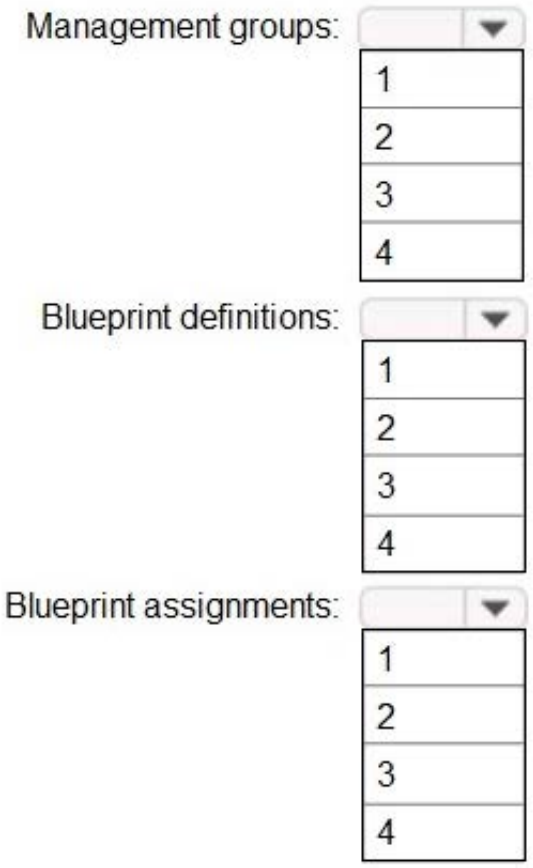

Correct Answer:

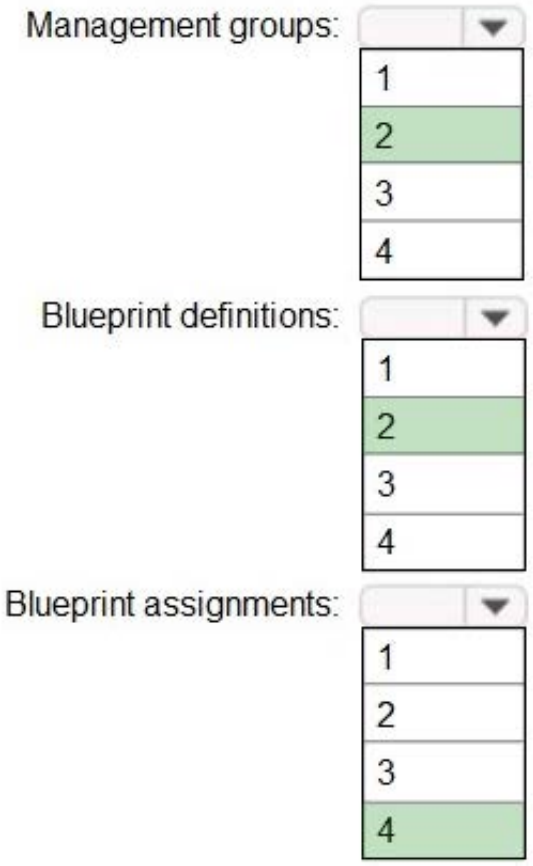

Box 1: 2

One management group for East, and one for West.

When creating a blueprint definition, you\\'ll define where the blueprint is saved. Blueprints can be saved to a management group or subscription that you have Contributor access to. If the location is a management group, the blueprint is

available to assign to any child subscription of that management group.

Box 2: 2

Box 3: 4

One assignment for each subscription.

"Assigning a blueprint definition to a management group means the assignment object exists at the management group. The deployment of artifacts still targets a subscription. To perform a management group assignment, the Create Or

Update REST API must be used and the request body must include a value for properties.scope to define the target subscription." https://docs.microsoft.com/en-us/azure/governance/blueprints/overview#blueprint-assignment

QUESTION 7

HOTSPOT

You have an Azure subscription.

You create a storage account that will store documents.

You need to configure the storage account to meet the following requirements:

1.

Ensure that retention policies are standardized across the subscription.

2.

Ensure that data can be purged if the data is copied to an unauthorized location.

Which two settings should you enable? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

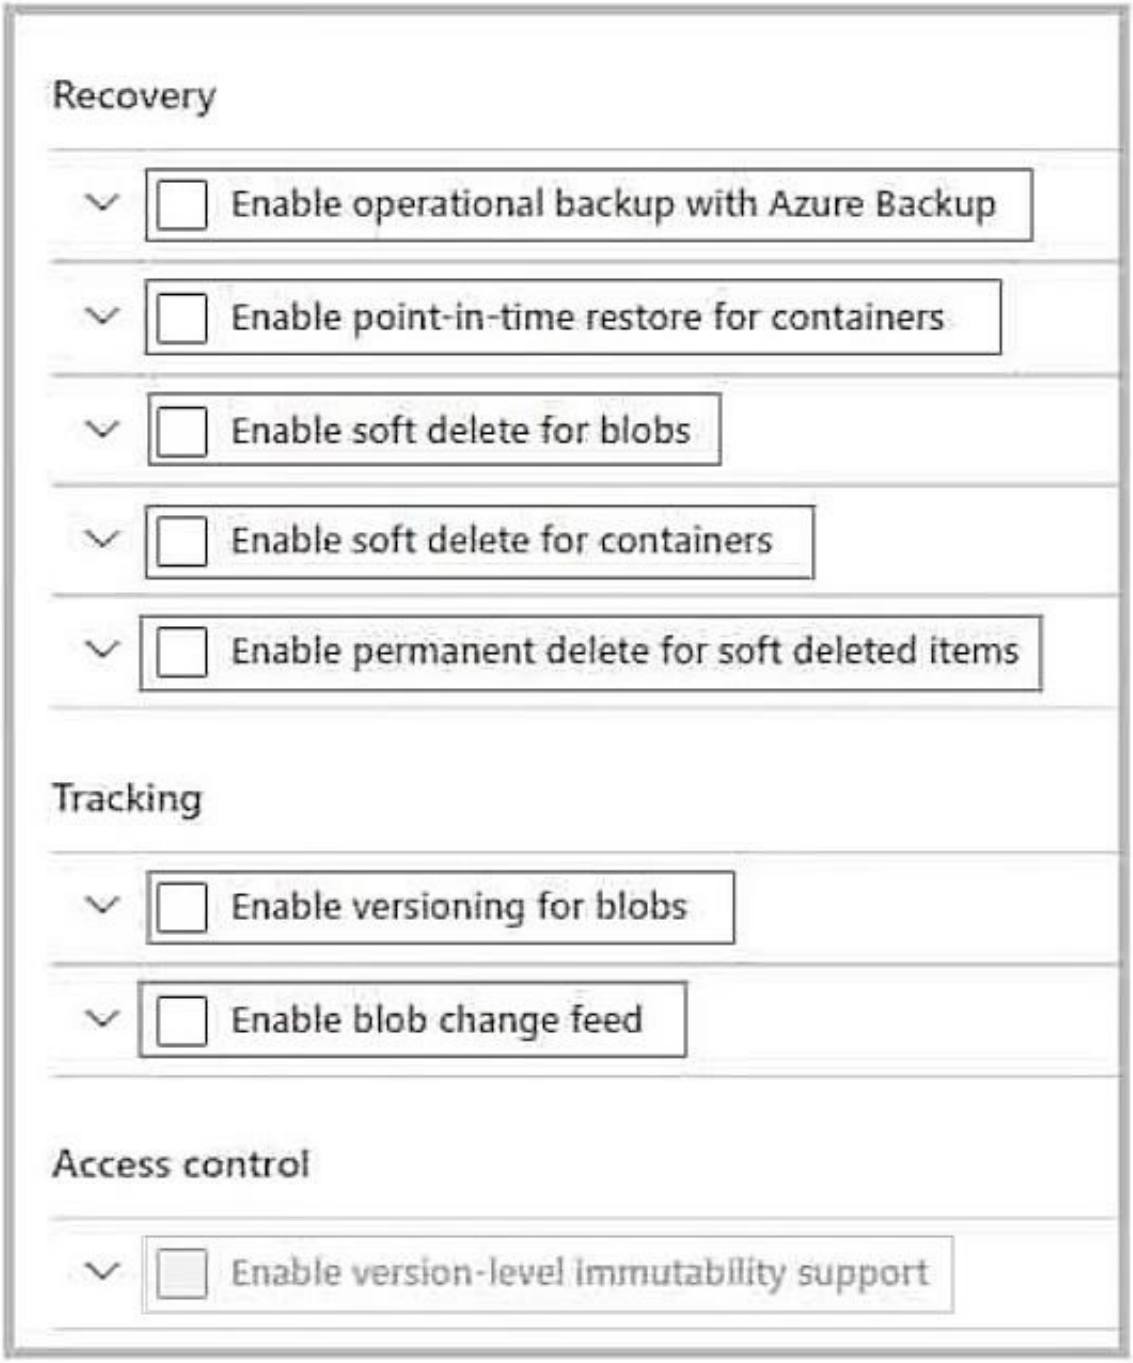

Correct Answer:

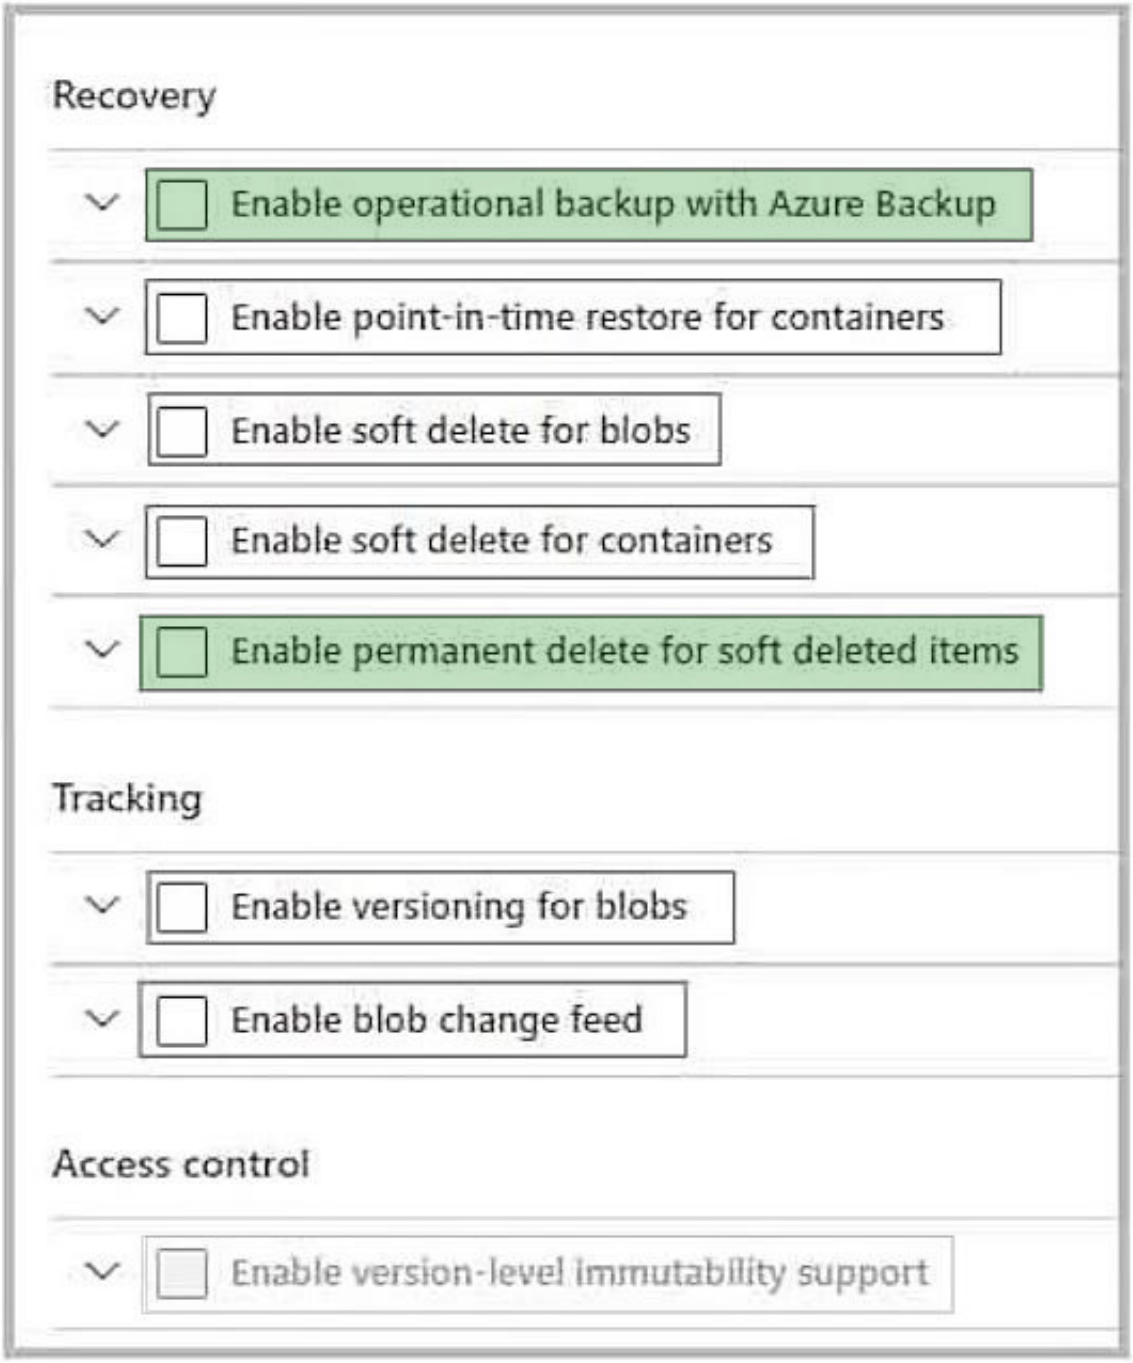

Box 1: Enable soft delete for containers

Ensure that retention policies are standardized across the subscription.

Container soft delete protects your data from being accidentally or erroneously modified or deleted. When container soft delete is enabled for a storage account, a container and its contents may be recovered after it has been deleted, within a

retention period that you specify.

Box 2: Enable permanent delete for soft deleted items

Ensure that data can be purged if the data is copied to an unauthorized location.

Incorrect:

*

Enable versioning for blobs You can enable Blob storage versioning to automatically maintain previous versions of a blob when it is modified or deleted. When blob versioning is enabled, then you can restore an earlier version of a blob to recover your data if it is erroneously modified or deleted.

*

Enable version-level immutability support Immutable storage for Azure Blob Storage enables users to store business-critical data in a WORM (Write Once, Read Many) state. While in a WORM state, data can\\'t be modified or deleted for a user-specified interval. By configuring immutability policies for blob data, you can protect your data from overwrites and deletes. Immutability policies include time-based retention policies and legal holds.

Reference: https://learn.microsoft.com/en-us/azure/storage/blobs/soft-delete-container-enable

QUESTION 8

HOTSPOT

You have an Azure subscription that contains an Azure key vault named KV1 and a virtual machine named VM1. VM1 runs Windows Server 2022: Azure Edition.

You plan to deploy an ASP.Net Core-based application named App1 to VM1.

You need to configure App1 to use a system-assigned managed identity to retrieve secrets from KV1. The solution must minimize development effort.

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

NOTE: Each correct selection is worth one point.

Hot Area:

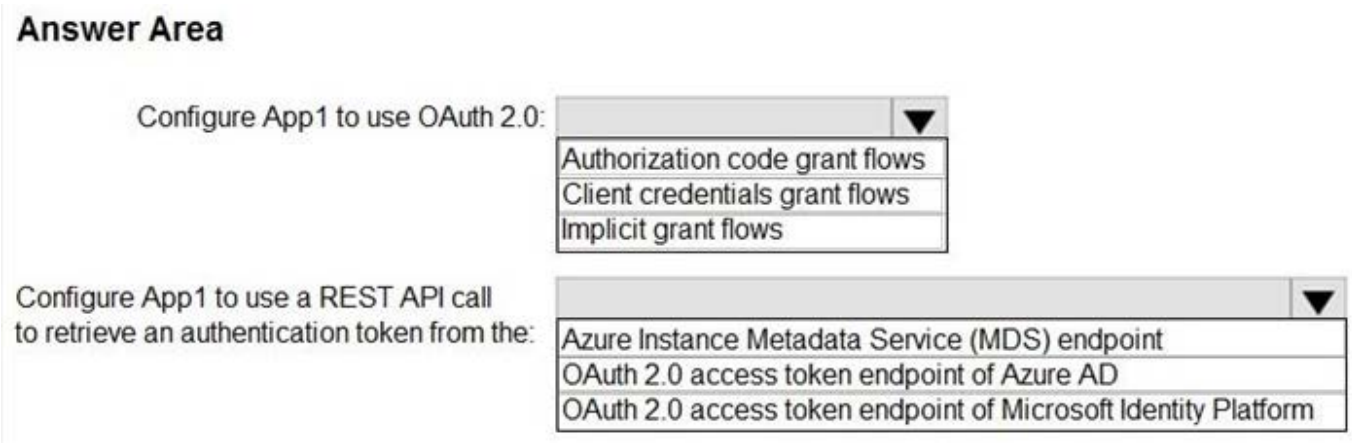

Correct Answer:

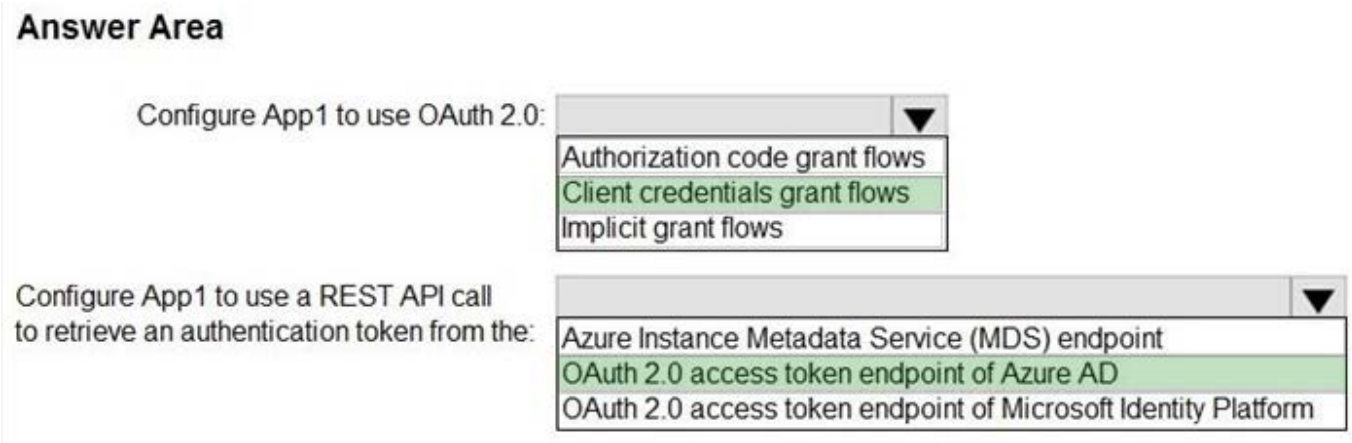

Box 1: Client Credentials flow Client Credentials flow - The only flow that does not require immediate user interaction, usually used when the OAuth client is acting on-behalf of itself, when user-consent doesn\\'t make sense, or when authorization primitives could be configured out-of-band (for instance via Azure AD)

Note: Authenticating to Azure Services Local machines don\\'t support managed identities for Azure resources. As a result, the Microsoft.Azure.Services.AppAuthentication library uses your developer credentials to run in your local development environment. When the solution is deployed to Azure, the library uses a managed identity to switch to an OAuth 2.0 client credential grant flow. This approach means you can test the same code locally and remotely without worry.

Incorrect:

*

Authorization code flow - Requires user interaction and consent, typically via the web browser, to get a code which is then used to issue an access token.

*

Implicit grant flow - Created for single page web / mobile webview apps, where token creation and handling is done entirely from the front end.

Box 2: OAuth 2.0 access token endpoint of Azure AD

Example: Issuing and inspecting our first OAuth token

At this stage, we should be able to issue tokens to Service A, on behalf of Service B - let\\'s see that in action.

In Azure AD application registration blade, go to Service B (as shown in previous steps)

In the Overview blade, Click on the ‘Endpoints

QUESTION 9

You store web access logs data in Azure Blob storage. You plan to generate monthly reports from the access logs.

You need to recommend an automated process to upload the data to Azure SQL Database every month.

What should you include in the recommendation?

A. Azure Data Factory

B. Data Migration Assistant

C. Microsoft SQL Server Migration Assistant (SSMA)

D. AzCopy

Correct Answer: A

Azure Data Factory is the platform that solves such data scenarios. It is the cloud-based ETL and data integration service that allows you to create data-driven workflows for orchestrating data movement and transforming data at scale. Using Azure Data Factory, you can create and schedule data-driven workflows (called pipelines) that can ingest data from disparate data stores. You can build complex ETL processes that transform data visually with data flows or by using compute services such as Azure HDInsight Hadoop, Azure Databricks, and Azure SQL Database.

Reference: https://docs.microsoft.com/en-gb/azure/data-factory/introduction

QUESTION 10

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

Your company has deployed several virtual machines (VMs) on-premises and to Azure. Azure ExpressRoute has been deployed and configured for on-premises to Azure connectivity.

Several VMs are exhibiting network connectivity issues.

You need to analyze the network traffic to determine whether packets are being allowed or denied to the VMs.

Solution: Use the Azure Traffic Analytics solution in Azure Log Analytics to analyze the network traffic.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Instead use Azure Network Watcher to run IP flow verify to analyze the network traffic.

Reference: https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-monitoring-overview
https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-ip-flow-verify-overview

QUESTION 11

You Wave an Acme Directory forest named contoso.com. You install and configure Azure AD Connect to use password hath synchronization as the single sign- on (SSO) method Staging mode is enabled.

You review the synchronization results and discover that the Synchronization Service Manager does not display any sync jobs.

You need to ensure that the synchronization completes successfully.

What should you do?

A. From Synchronization Service Manager, run a full import

B. From Azure PowerShell, run Start-AdSyncCycle -PolicyType initial.

C. Run Azure AD Connect and set the SSO method to Pass-through Authentication

D. Run Azure AD Connect and disable staging mode._

Correct Answer: D

QUESTION 12

You plan to deploy an app that will use an Azure Storage account.

You need to deploy the storage account. The solution must meet the following requirements:

1.

Store the data of multiple users.

2.

Encrypt each user\\'s data by using a separate key.

3.

Encrypt all the data in the storage account by using Microsoft keys or customer-managed keys. What should you deploy?

A. files in a general purpose v2 storage account.

B. blobs in an Azure Data Lake Storage Gen2 account.

C. files in a premium file share storage account.

D. blobs in a general purpose v2 storage account

Correct Answer: B

QUESTION 13

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:

1.

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.

2.

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

Correct Answer: D

When you create a function app in Azure, you must choose a hosting plan for your app. There are three basic hosting plans available for Azure Functions: Consumption plan, Premium plan, and Dedicated (App Service) plan. For the Consumption plan, you don\\'t have to pay for idle VMs or reserve capacity in advance.

Connect to private endpoints with Azure Functions As enterprises continue to adopt serverless (and Platform-as-a-Service, or PaaS) solutions, they often need a way to integrate with existing resources on a virtual network. These existing resources could be databases, file storage, message queues or event streams, or REST APIs.

Reference: https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale
https://techcommunity.microsoft.com/t5/azure-functions/connect-to-private-endpoints-with-azure-functions/ba-p/1426615
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#hosting-plans-comparison

QUESTION 14

HOTSPOT

You have Azure subscription that contains a virtual network named VNet1. VNet1 includes a VPN gateway, uses an IP address space of 10.0.0.0/16 and contains the subnets in the following table.

Name	IP address range
Subnet0	10.0.0.0/24
Subnet1	10.0.1.0/24
Subnet2	10.0.2.0/24
GatewaySubnet	10.0.254.0/24

Subnet1 contains a virtual appliance named VM1 that operates as a router.

You create a routing table named RT1.

You need to route all inbound traffic from the VPN gateway to VNet1 through VM1.

How should you configure RT1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

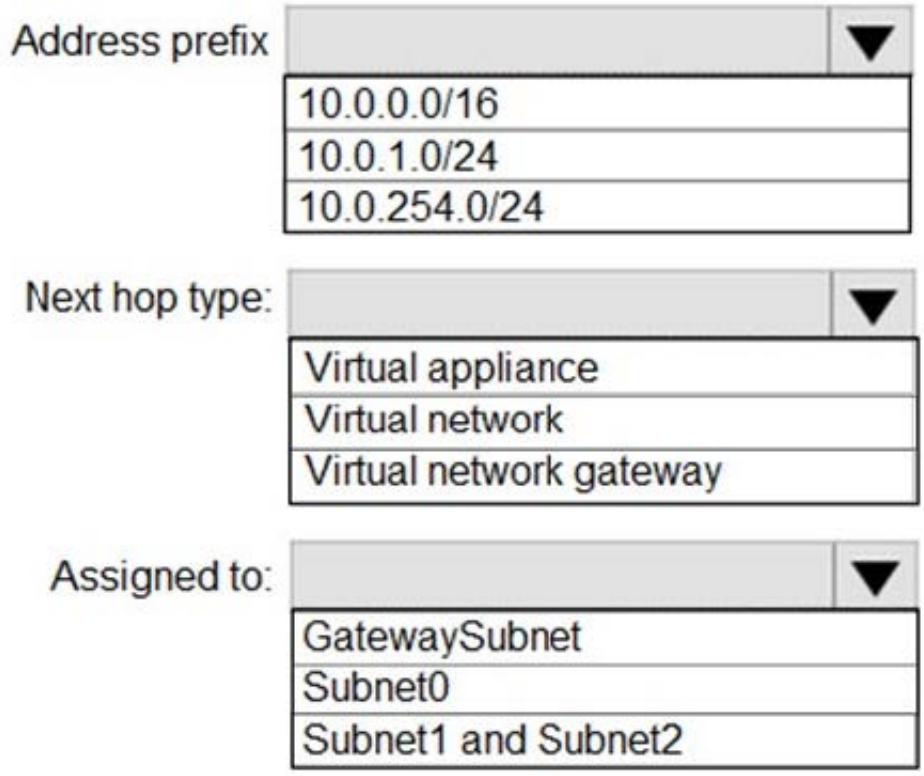

Correct Answer:

Answer Area

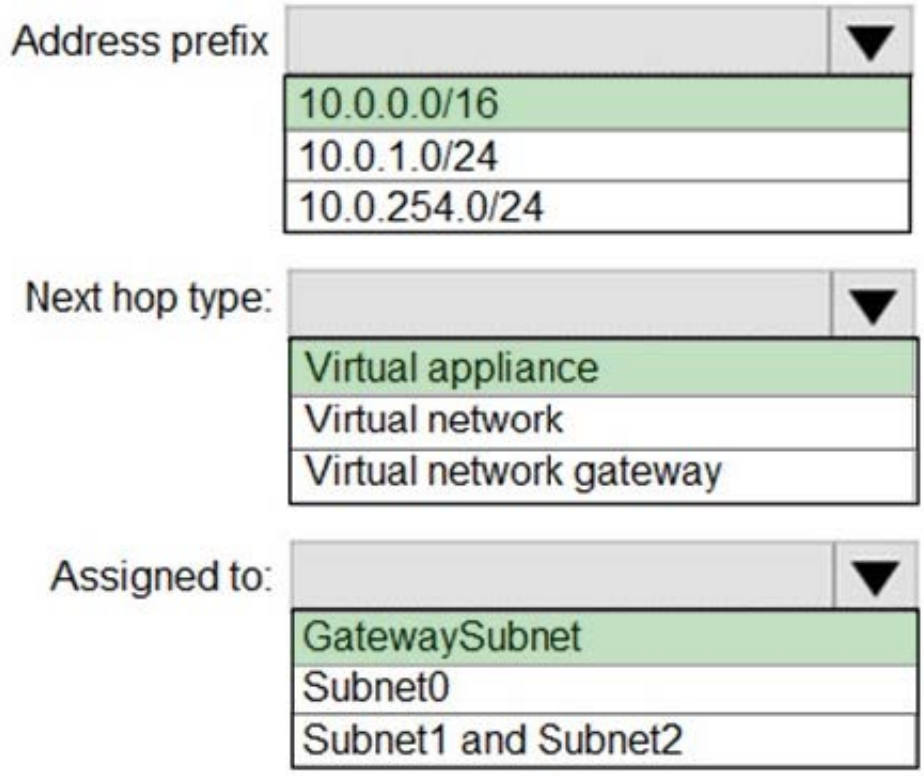

QUESTION 15

HOTSPOT

You have two on-premises Microsoft SQL Server 2017 instances that host an Always On availability group named AG1. AG1 contains a single database named DB1.

You have an Azure subscription that contains a virtual machine named VM1VM1 runs Linux and contains a SQL Server 2019 instance.

You need to migrate DB1 to VMI. The solution must minimize downtime on DBI. What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

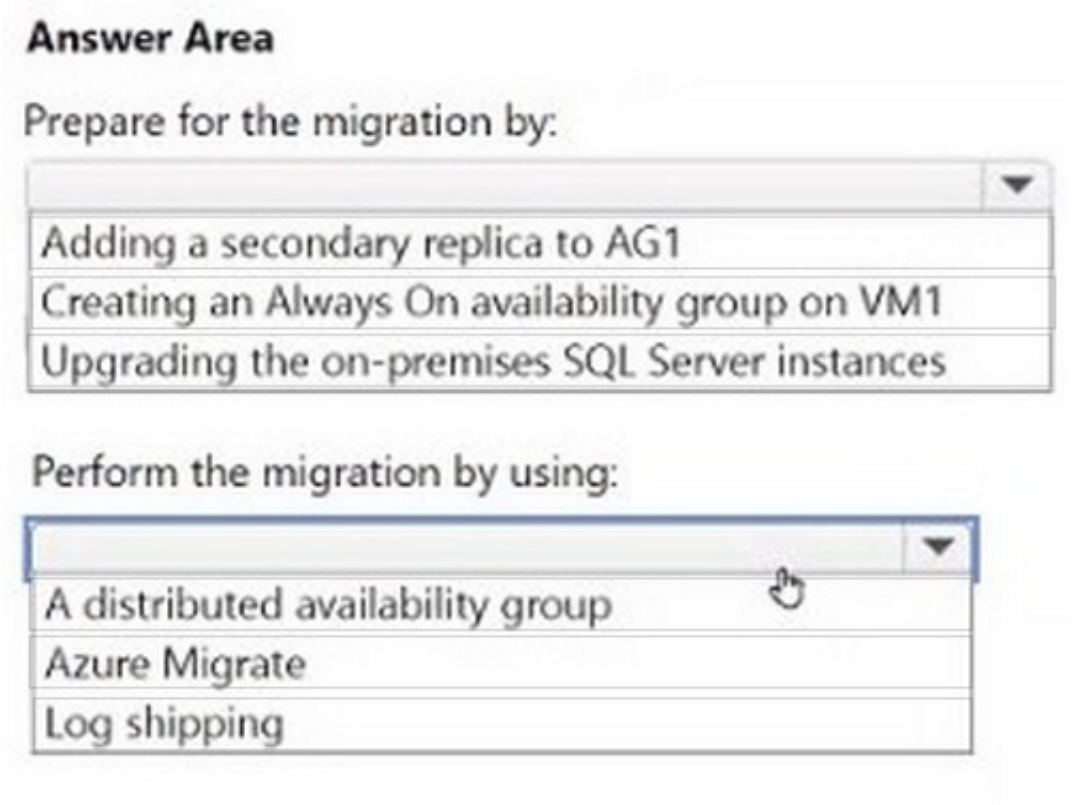

Correct Answer: