

Latest Version: 22.0

Question: 1

HOTSPOT

You need to ensure that Azure Data Factory pipelines can be deployed. How should you configure authentication and authorization for deployments? To answer, select the appropriate options in the answer choices.

NOTE: Each correct selection is worth one point.

Security requirement

Authorization

Authentication

Technology

RBAC	v
DAC	
MAC	
Claims	

Service Principal	^
Kerberos	
Certificate-based	
Bearer Token	

Answer:

Security requirement

Technology

Authorization

RBAC	v
DAC	
MAC	
Claims	

Authentication

Service Principal	^
Kerberos	
Certificate-based	
Bearer Token	

Explanation:

The way you control access to resources using RBAC is to create role assignments. This is a key concept to understand – it's how permissions are enforced. A role assignment consists of three elements: security principal, role definition, and scope.

Scenario:

No credentials or secrets should be used during deployments

Phone-based poll data must only be uploaded by authorized users from authorized devices

Contractors must not have access to any polling data other than their own

Access to polling data must set on a per-active directory user basis

References:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/overview>

Question: 2

DRAG DROP

You need to provision the polling data storage account.

How should you configure the storage account? To answer, drag the appropriate Configuration Value to the correct Setting. Each Configuration Value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Configuration values**Answer Area****Setting****Configuration value**

Account type

Replication type

Answer:**Configuration values****Answer Area****Setting****Configuration value**

Account type

Replication type

Explanation:

Account type: StorageV2

You must create new storage accounts as type StorageV2 (general-purpose V2) to take advantage of Data Lake Storage Gen2 features.

Scenario: Polling data is stored in one of the two locations:

An on-premises Microsoft SQL Server 2019 database named PollingData

Azure Data Lake Gen 2

Data in Data Lake is queried by using PolyBase

Replication type: RA-GRS

Scenario: All services and processes must be resilient to a regional Azure outage.

Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from the primary region. If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable.

If you opt for GRS, you have two related options to choose from:

GRS replicates your data to another data center in a secondary region, but that data is available to be read only if Microsoft initiates a failover from the primary to secondary region.

Read-access geo-redundant storage (RA-GRS) is based on GRS. RA-GRS replicates your data to another data center in a secondary region, and also provides you with the option to read from the secondary region. With RA-GRS, you can read from the secondary region regardless of whether Microsoft initiates a failover from the primary to secondary region.

References:

<https://docs.microsoft.com/bs-cyrl-ba/azure/storage/blobs/data-lake-storage-quickstart-create-account>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs>

Question: 3

HOTSPOT

You need to ensure polling data security requirements are met.

Which security technologies should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Context	Security technology
SQL Server	Azure Active Directory user
	Domain Active Directory user
	Managed Identity
PolyBase	Database scoped credential
	Database encryption key
	Application role

Answer:

Context Security technology

SQL Server

Azure Active Directory user	v
Domain Active Directory user	
Managed Identity	

PolyBase

Database scoped credential	v
Database encryption key	
Application role	

Explanation:

Box 1: Azure Active Directory user

Scenario:

Access to polling data must set on a per-active directory user basis

Box 2: DataBase Scoped Credential

SQL Server uses a database scoped credential to access non-public Azure blob storage or Kerberos-secured Hadoop clusters with PolyBase.

PolyBase cannot authenticate by using Azure AD authentication.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-scoped-credential-transact-sql>

Question: 4

DRAG DROP

You need to ensure that phone-based polling data can be analyzed in the PollingData database.

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

Actions

Parameterize deployment by using Azure Integration Runtime

Configure an Azure Logic App to deploy the deployment artifact

Configure Azure DevOps to deploy the deployment artifact

Create a deployment artifact containing an extracted Azure Resource Manager template

Parameterize deployment by using the Azure Resource Manager template parameter file

Create a deployment artifact containing a SQL Server Integration Services (SSIS) package

Answer Area

Answer:

Answer Area

Create a deployment artifact containing an extracted Azure Resource Manager template

Parameterize deployment by using the Azure Resource Manager template parameter file

Configure Azure DevOps to deploy the deployment artifact

Explanation:

Scenario:

All deployments must be performed by using Azure DevOps. Deployments must use templates used in multiple environments

No credentials or secrets should be used during deployments

Question: 5

You need to ensure that phone-based polling data can be analyzed in the PollingData database. How should you configure Azure Data Factory?

- A. Use a tumbling schedule trigger
- B. Use an event-based trigger
- C. Use a schedule trigger
- D. Use manual execution

Answer: C

Explanation:

When creating a schedule trigger, you specify a schedule (start date, recurrence, end date etc.) for the trigger, and associate with a Data Factory pipeline.