

Testpassport**Q&A**



H i g h e r Q u a l i t y

B e t t e r S e r v i c e !

We offer free update service for one year
[Http://www.testpassport.com](http://www.testpassport.com)

Exam : **DP-200**

Title : **Implementing an Azure Data
Solution**

Version : **DEMO**

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

Actions

Answer Area

| | |
|--|--|
| Provision an Azure SQL Data Warehouse instance. | |
| Connect to the Blob storage container by using SQL Server Management Studio. | |
| Provision an Azure Blob storage container. | |
| Run Transact-SQL statements to load data. | |
| Connect to the Azure SQL Data Warehouse by using SQL Server Management Studio. | |
| Build external tables by using Azure portal. | |
| Build external tables by using SQL Server Management Studio. | |

Answer:

Actions**Answer Area**

| | |
|--|--|
| Provision an Azure SQL Data Warehouse instance. | Provision an Azure SQL Data Warehouse instance. |
| Connect to the Blob storage container by using SQL Server Management Studio. | Connect to the Blob storage container by using SQL Server Management Studio. |
| Provision an Azure Blob storage container. | Build external tables by using SQL Server Management Studio. |
| Run Transact-SQL statements to load data. | Run Transact-SQL statements to load data. |
| Connect to the Azure SQL Data Warehouse by using SQL Server Management Studio. | |
| Build external tables by using Azure portal. | |
| Build external tables by using SQL Server Management Studio. | |

Explanation:

Step 1: Provision an Azure SQL Data Warehouse instance.

Create a data warehouse in the Azure portal.

Step 2: Connect to the Azure SQL Data warehouse by using SQL Server Management Studio

Connect to the data warehouse with SSMS (SQL Server Management Studio)

Step 3: Build external tables by using the SQL Server Management Studio

Create external tables for data in Azure blob storage.

You are ready to begin the process of loading data into your new data warehouse. You use external tables to load data from the Azure storage blob.

Step 4: Run Transact-SQL statements to load data.

You can use the CREATE TABLE AS SELECT (CTAS) T-SQL statement to load the data from Azure Storage Blob into new tables in your data warehouse.

References:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/sql-data-warehouse/load-data-from-azure-blob-storage-using-polybase.md>

2.You develop data engineering solutions for a company. The company has on-premises Microsoft SQL Server databases at multiple locations.

The company must integrate data with Microsoft Power BI and Microsoft Azure Logic Apps. The solution must avoid single points of failure during connection and transfer to the cloud. The solution must also minimize latency. You need to secure the transfer of data between on-premises databases and Microsoft Azure.

What should you do?

- A. Install a standalone on-premises Azure data gateway at each location
- B. Install an on-premises data gateway in personal mode at each location
- C. Install an Azure on-premises data gateway at the primary location
- D. Install an Azure on-premises data gateway as a cluster at each location

Answer: D

Explanation:

You can create high availability clusters of On-premises data gateway installations, to ensure your organization can access on-premises data resources used in Power BI reports and dashboards. Such clusters allow gateway administrators to group gateways to avoid single points of failure in accessing on-premises data resources. The Power BI service always uses the primary gateway in the cluster, unless it's not available. In that case, the service switches to the next gateway in the cluster, and so on.

References: <https://docs.microsoft.com/en-us/power-bi/service-gateway-high-availability-clusters>

3.You are a data architect. The data engineering team needs to configure a synchronization of data between an on-premises Microsoft SQL Server database to Azure SQL Database.

Ad-hoc and reporting queries are being overutilized the on-premises production instance.

The synchronization process must:

- Perform an initial data synchronization to Azure SQL Database with minimal downtime
- Perform bi-directional data synchronization after initial synchronization

You need to implement this synchronization solution.

Which synchronization method should you use?

- A. transactional replication
- B. Data Migration Assistant (DMA)
- C. backup and restore
- D. SQL Server Agent job
- E. Azure SQL Data Sync

Answer: E

Explanation:

SQL Data Sync is a service built on Azure SQL Database that lets you synchronize the data you select bidirectionally across multiple SQL databases and SQL Server instances.

With Data Sync, you can keep data synchronized between your on-premises databases and Azure SQL databases to enable hybrid applications.

Compare Data Sync with Transactional Replication

| | Data Sync | Transactional Replication |
|---------------|--|---|
| Advantages | <ul style="list-style-type: none"> - Active-active support - Bi-directional between on-premises and Azure SQL Database | <ul style="list-style-type: none"> - Lower latency - Transactional consistency - Reuse existing topology after migration |
| Disadvantages | <ul style="list-style-type: none"> - 5 min or more latency - No transactional consistency - Higher performance impact | <ul style="list-style-type: none"> - Can't publish from Azure SQL Database single database or pooled database - High maintenance cost |

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-sync-data>

4. An application will use Microsoft Azure Cosmos DB as its data solution. The application will use the Cassandra API to support a column-based database type that uses containers to store items.

You need to provision Azure Cosmos DB.

Which container name and item name should you use? Each correct answer presents part of the solutions. NOTE: Each correct answer selection is worth one point.

- A. collection
- B. rows
- C. graph
- D. entities
- E. table

Answer: B,E

Explanation:

B: Depending on the choice of the API, an Azure Cosmos item can represent either a document in a collection, a row in a table or a node/edge in a graph.

The following table shows the mapping between API-specific entities to an Azure Cosmos item:

| Cosmos entity | SQL API | Cassandra API | Azure Cosmos DB's API for MongoDB | Gremlin API | Table API |
|-------------------|----------|---------------|-----------------------------------|--------------|-----------|
| Azure Cosmos item | Document | Row | Document | Node or Edge | Item |

E: An Azure Cosmos container is specialized into API-specific entities as follows:

| Azure Cosmos entity | SQL API | Cassandra API | Azure Cosmos DB's API for MongoDB | Gremlin API | Table API |
|------------------------|------------|---------------|-----------------------------------|-------------|-----------|
| Azure Cosmos container | Collection | Table | Collection | Graph | Table |

References: <https://docs.microsoft.com/en-us/azure/cosmos-db/databases-containers-items>

5. A company has a SaaS solution that uses Azure SQL Database with elastic pools. The solution contains a dedicated database for each customer organization. Customer organizations have peak usage at different periods during the year. You need to implement the Azure SQL Database elastic pool to minimize cost.

Which option or options should you configure?

- A. Number of transactions only
- B. eDTUs per database only
- C. Number of databases only
- D. CPU usage only
- E. eDTUs and max data size

Answer: E

Explanation:

The best size for a pool depends on the aggregate resources needed for all databases in the pool. This involves determining the following:

Maximum resources utilized by all databases in the pool (either maximum DTUs or maximum vCores depending on your choice of resourcing model).

Maximum storage bytes utilized by all databases in the pool.

Note: Elastic pools enable the developer to purchase resources for a pool shared by multiple databases to accommodate unpredictable periods of usage by individual databases. You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>