You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1.

You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1.

Solution: You use an Azure Synapse Analytics serverless SQL pool to create an external table that has an additional DateTime column.

Does this meet the goal?

A. Yes
B. No

Correct Answer: A
Explanation:
In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables

QUESTION 66
You have an Azure SQL managed instance named SQL1 and two Azure web apps named App1 and App2.

You need to limit the number of IOPs that App2 queries generate on SQL1.

Which two actions should you perform on SQL1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Enable query optimizer fixes.
B. Enable Resource Governor.
C. Enable parameter sniffing.
D. Create a workload group.
E. Configure In-memory OLTP.
F. Run the Database Engine Tuning Advisor.
G. Reduce the Max Degree of Parallelism value.

Correct Answer: BC
Explanation:
https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15

QUESTION 67
You have an Azure SQL database named db1 on a server named server1.

You need to modify the MAXDOP settings for db1.

What should you do?

A. Connect to db1 and run the sp_configure command.
B. Connect to the master database of server1 and run the sp_configure command.
C. Configure the extended properties of db1.
D. Modify the database scoped configuration of db1.

Correct Answer: D
Explanation:
https://docs.microsoft.com/en-us/azure/azure-sql/database/configure-max-degree-of-parallelism

QUESTION 68
You have an Azure SQL database that contains a table named factSales. FactSales contains the columns shown in the following table.

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(30)

FactSales has 6 billion rows and is loaded nightly by using a batch process.

Which type of compression provides the greatest space reduction for the database?

A. page compression
B. row compression
C. columnstore compression
D. columnstore archival compression

Correct Answer: D
Explanation:
Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.

Note:
Columnstore -- The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.

Reference:
https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression

QUESTION 69
Your on-premises network contains a Microsoft SQL Server 2016 server that hosts a database

named db1.

You have an Azure subscription.

You plan to migrate db1 to an Azure SQL managed instance.

You need to create the SQL managed instance. The solution must minimize the disk latency of the instance.

Which service tier should you use?

A. Hyperscale
B. General Purpose
C. Premium
D. Business Critical

Correct Answer: A

QUESTION 70
You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in fiveminute intervals and report only events that arrive during the interval.

The output will be sent to a Delta Lake table.

Which output mode should you use?

A. complete
B. append
C. update

Correct Answer: A
Explanation:
Complete mode: You can use Structured Streaming to replace the entire table with every batch.

Reference:
https://docs.databricks.com/delta/delta-streaming.html

QUESTION 71
You are designing a date dimension table in an Azure Synapse Analytics dedicated SQL pool. The date dimension table will be used by all the fact tables. Which distribution type should you recommend to minimize data movement?

A. HASH
B. REPLICATE
C. ROUND_ROBIN
D. None of the above

Correct Answer: B
Explanation:
A replicated table has a full copy of the table available on every Compute node. Queries run fast on replicated tables since joins on replicated tables don't require data movement. Replication requires extra storage, though, and isn't practical for large tables.

Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribute

QUESTION 72
You have a new Azure subscription.

You create an Azure SQL Database instance named DB1 on an Azure SQL Database server named Server1.

You need to ensure that users can connect to DB1 in the event of an Azure regional outage. In the event of an outage, applications that connect to DB1 must be able to connect without having to update the connection strings.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. From the properties of DB1. configure geo-replication.
B. From the properties of Server1 add a failover group.
C. Create a new Azure SQL Database server named Server2.
D. From the properties of Server1 configure retention for DB1
E. Create a new Azure SQL Database instance named DB2.

Correct Answer: BC
Explanation:
https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview?tabs=azure-powershell#best-practices-for-sql-database
https://docs.microsoft.com/en-us/azure/azure-sql/database/failover-group-add-single-database-tutorial?tabs=azure-portal

QUESTION 73
Your on-premises network contains a server that hosts a 60-TB database named DB 1. The network has a 10-Mbps internet connection.

You need to migrate DB 1 to Azure. The solution must minimize how long it takes to migrate the database.

What should you use?

A. Azure Migrate
B. Data Migration Assistant (DMA)
C. Azure Data BOX

D.   Azure Database Migration Service

Correct Answer: D
Explanation:
https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-with-Microsoft-Azure-Migrate

QUESTION 74
HOTSPOT
You configure version control for an Azure Data Factory instance as shown in the following exhibit.

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.