SQL
KONFERENZ
2025



What is new in Fabric DW

Jovan Popovic

SQL KONFERENZ 2025

Agenda

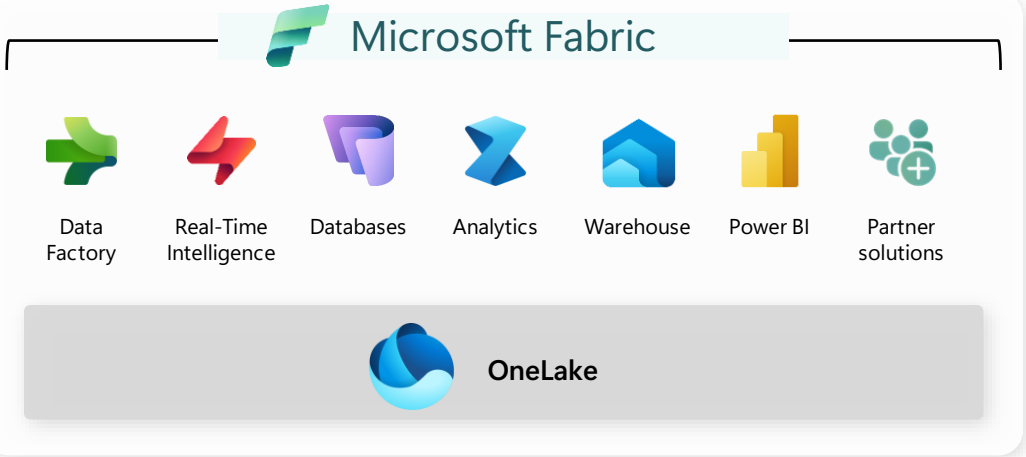
- What is Fabric DW
- New features in Fabric DW



SQL KONFERENZ 2025

What is Fabric Datawarehouse?





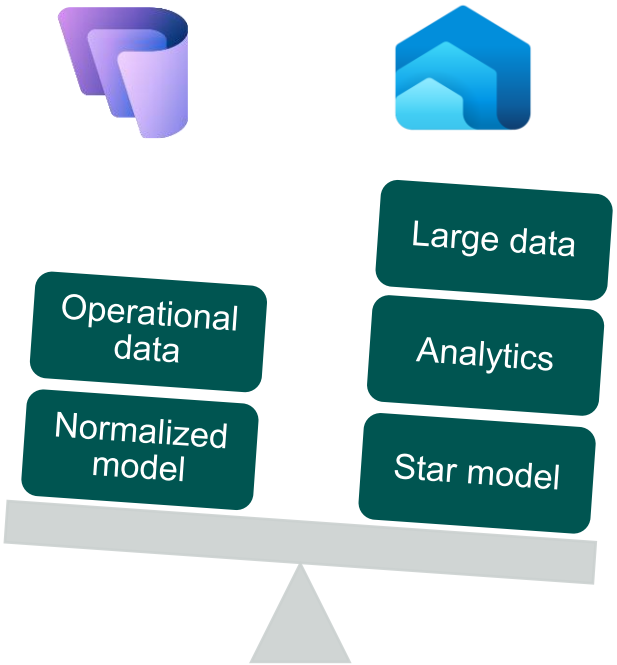
SQL KONFERENZ 2025

What is Fabric Datawarehouse?



- Modern SaaS warehouse
- Fully integrated into Fabric ecosystem
- Multi-PB analytic workloads
- Open and extensible
- Elasticity and scalability

Database vs warehouse





SQL
KONFERENZ
2025



New Functionalities

Why is Fabric Dw the best choice for analytics?

Recently added features

- Open mirroring
- Performance insights
- Query plans
- Query hints

- LOB/JSON/Spatial
- Spatial
- OPENROWSET
- Improved data loading

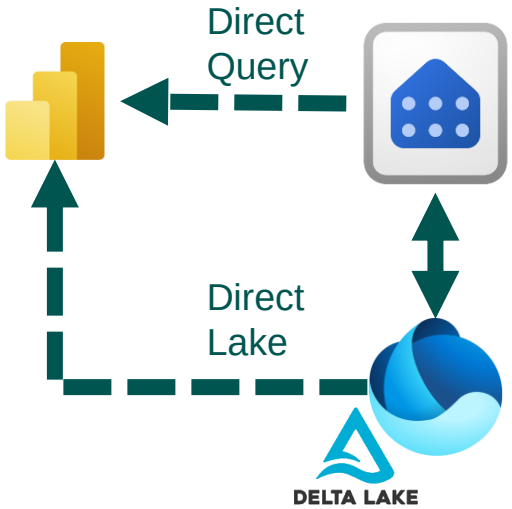


SQL KONFERENZ 2025

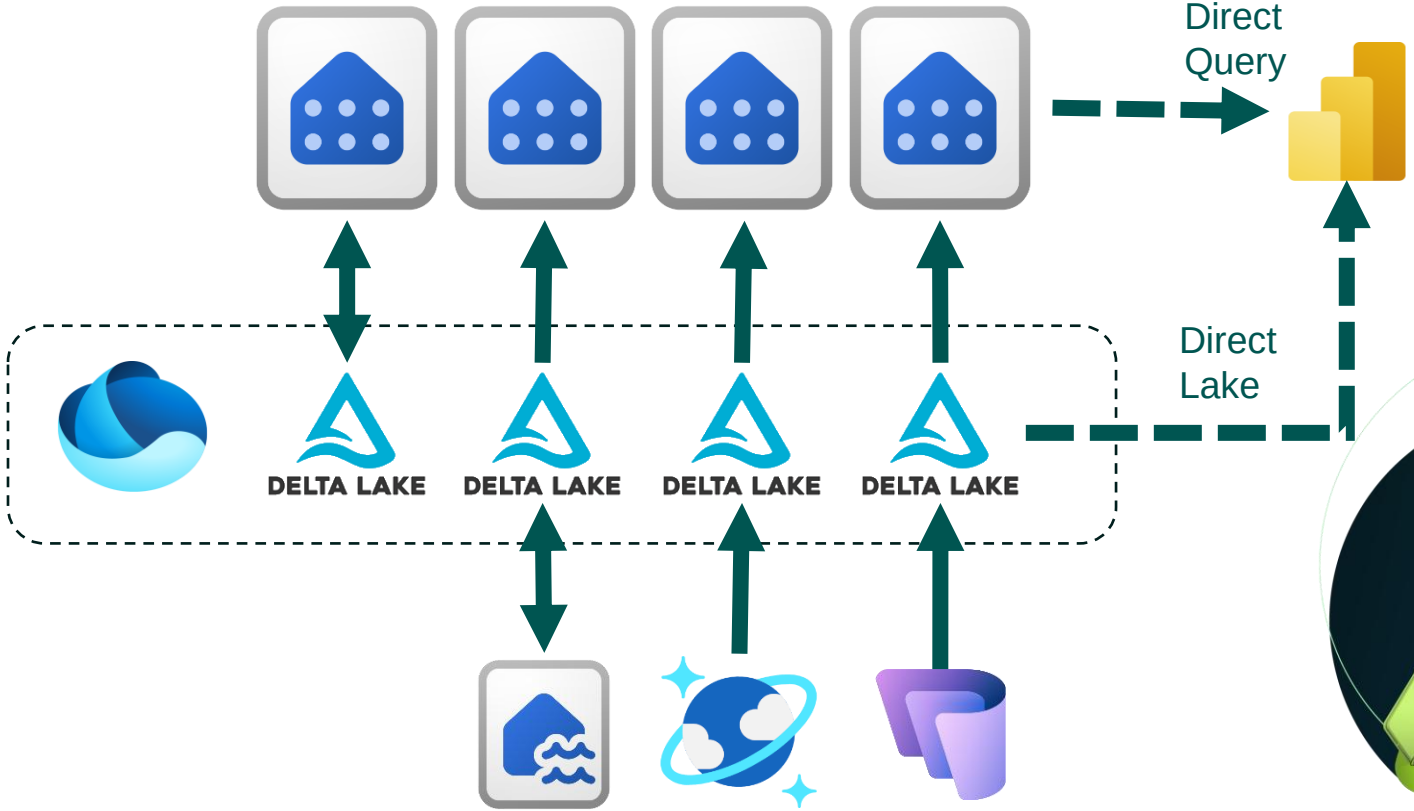
Separation of compute and storage

- Data is in One Lake
- Stored as Delta Lake format
- Enables dual data access
- Best fit for PowerBI



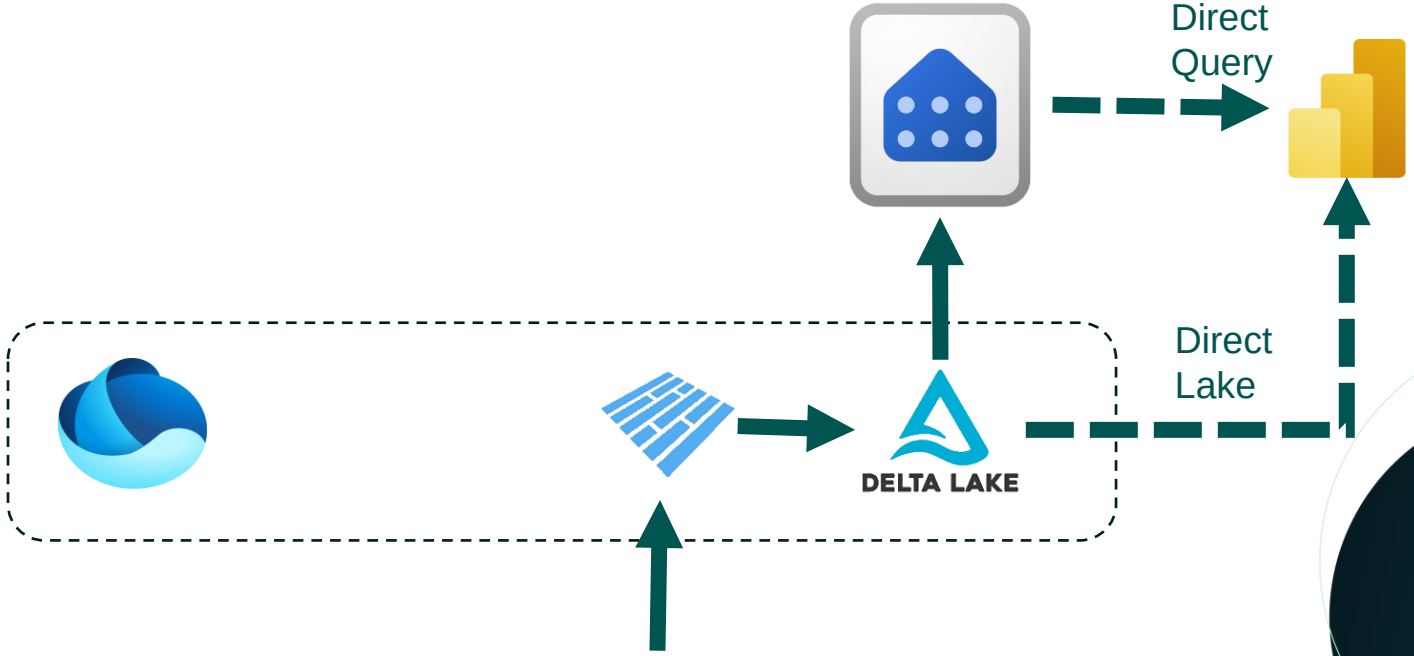


Analytics endpoint for any data source



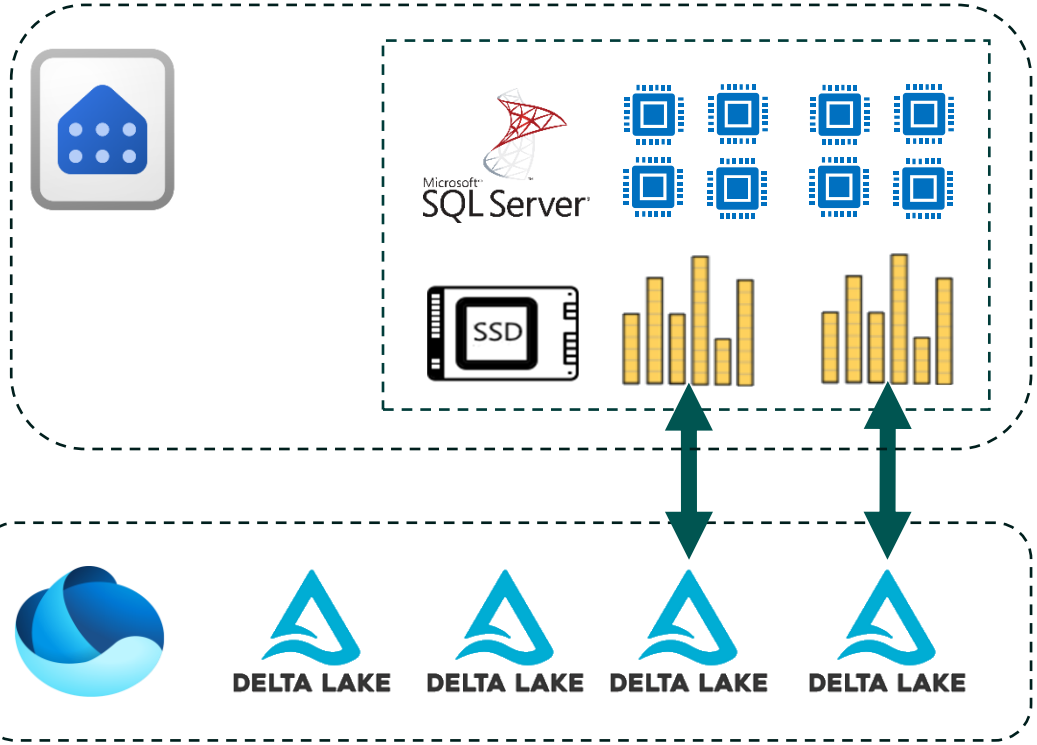


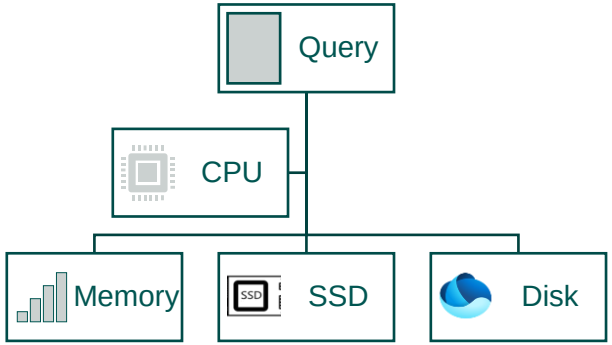
Open mirroring (preview)





Powered by SQL Server engine







Query insights execution stats (NEW)

- New columns are added in the queryinsights view

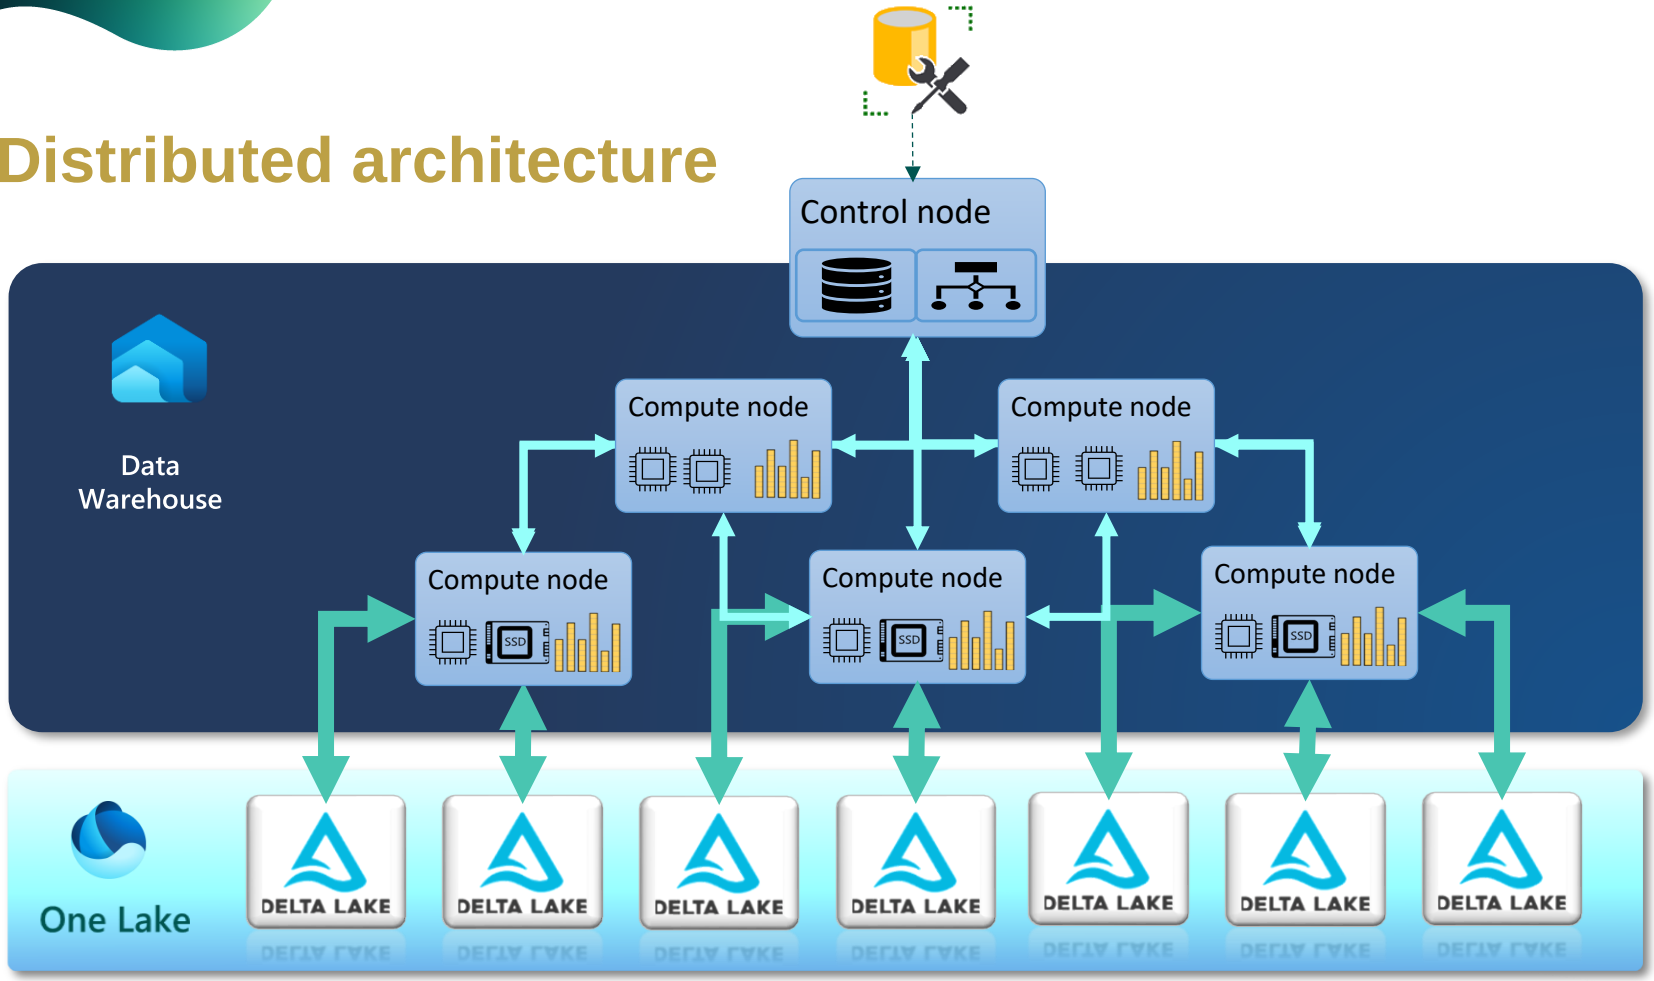
```
SELECT label, DATEDIFF(second,start_time, end_time) as duration_s,
       allocated_cpu_time_ms, data_scanned_remote_storage_mb,
       data_scanned_memory_mb, data_scanned_disk_mb
FROM queryinsights.exec_requests_history
WHERE label in ('WK-001-en','WK-002-en','WK-003-en','WK-004-en')
AND status = 'Succeeded'
```

150 %

Results | Messages

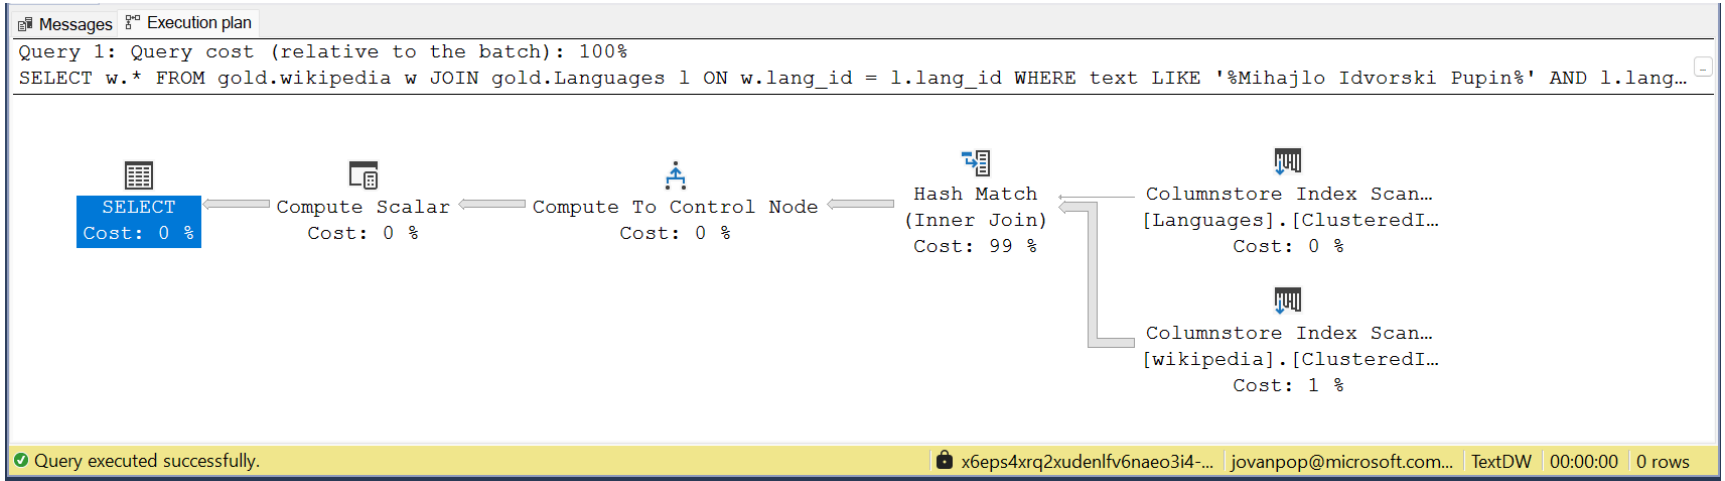
	label	duration_s	allocated_cpu_time_ms	data_scanned_remote_storage_mb	data_scanned_memory_mb	data_scanned_disk_mb
6	WK-002-en	4	184442	7236.545	52854.487	0.000
7	WK-002-en	2	60221	0.000	60091.032	0.000
8	WK-002-en	3	78483	0.000	60091.032	0.000

Distributed architecture



Estimated execution plan (NEW)

- Displaying estimated execution plan in tools
- SET SHOWPLAN_XML ON/OFF
- Includes information about node data transfer
- Does not show actual execution plan



LOB support

- VARCHAR(MAX)/VARBINARY(MAX) support up to 1MB per cell
  - UTF8 collations
- Improved JSON support
  - Public preview – JSON_OBJECTAGG, JSON_ARRAYAGG
- Spatial functions
  - Enabled spatial functions WKB/WKT conversions
  - Not enabled – spatial types and indexes

OPENROWSET/BULK INSERT

- OPENROWSET
  - Browse files in storage
  - Read partitioned data
  - Read Geoparquet
  - ETL
  - Incremental loads
- BULK INSERT
  - SQL Server syntax
  - FORMATFILE
  - DATAFILETYPE/CODEPAGE



OPENROWSET

- Read Parquet/CSV files on ADLS (No OneLake)
- Reference custom partitions
- Schema inference/WITH schema

```
SELECT  puYear = ors.filepath(1),
        puMonth = ors.filepath(2),
        ors.*
FROM OPENROWSET(
    BULK 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/puYear=*/puMonth=*/*.parquet',
    ROWS_PER_BATCH=1571671152)
    WITH (
            vendorId VARCHAR(100),
            pickupDateTime DATETIME2,
            dropoffDateTime DATETIME2,
            passengerCount INT,
            tripDistance FLOAT
    )  AS ors
WHERE ors.filepath(1) IN ('2001', '2002')
```

OPENROWSET limitations

	Supported	Not supported
File formats	Parquet, CSV	Delta, CosmosDB
Storage	Azure Blob storage, ADLS	OneLake, OneLake shortcuts
Authentication	EntraID passthrough, public storage	SAS/SAK, SPN, Workspace Identity, Managed Identity

BULK INSERT (preview)

- Enabling migration without code changes
- Based on COPY INTO
- Additional features: DATAFILETYPE, CODEPAGE, FORMATFILE
- Not supported
  - ERRORFILE
  - Datasources
  - Options - KEEPIDENTITY, FIRE_TRIGGERS, CHECK_CONSTRAINTS, TABLOCK, ORDER, ROWS_PER_BATCH, KILOBYTES_PER_BATCH, and BATCHSIZE

Conclusion

- What's new and planned for Data Warehouse in Microsoft Fabric - Microsoft Fabric | Microsoft Learn
  - Temp tables, MERGE, scalar functions, EXECUTE AS, result set caching
- Fabric Ideas – Data Warehouse
- Call to action – provide feedback!

Thank you!

Enjoy SQL Konferenz 2025!



SQL KONFERENZ 2025