



Context

... from the SQL-based data storage and processing perspective



Lakehouse & Warehouse

Lakehouse

Flexible data platform supporting structured, semi-structured, and unstructured data.

Supports notebooks, Spark, and SQL—ideal for advanced data workflows.

OLAP-focused with some OLTP-like capabilities, especially via Spark or Delta Lake support (which allows ACID transactions) Warehouse

SQL-optimized engine tailored for structured data analytics.

Dedicated, SQL-first engine designed for structured analytics at scale.

Strictly OLAP-focused. It's designed for structured analytics, reporting, and BI workloads that involve complex queries over large volumes of historical data.



Demo 01 – SQL Endpoints Look the same, but they are not ! ≌

Feature	Lakehouse Endpoint	Warehouse	
CREATE TABLE	×		
ALTER/DROP TABLE	×		
INSERT/UPDATE/DELETE	×		
CTAS (CREATE AS SELECT)	×		
Identity Columns	×	🗙 (Workaround)	
Temp Tables / CTEs	Views only		
Data Types	Mirrors Delta Lake	Subset T-SQL	
Read-only vs Read-write	Read-only	Read-write	
Endpoint availability	Auto for each LH	1 warehouse/item	

Something seemed to be missing in Fabric!

What Fabric Lakehouse and Warehouse don't provide :

▲ No native support for **SQL Server-style transactional logic** like stored procedures, triggers, or complex joins in real-time.

Limited suitability for line-of-business apps, which rely on OLTP models.

😵 The SQL Legacy Challenge ...

Decades of mission-critical SQL apps with deeply embedded business logic.

Extensive use of stored procedures, views, and functions

b High cost and complexity in rewriting SQL logic into Spark or notebooks.

SQL database in Microsoft Fabric

... The home in Fabric for OLTP workloads





Fabric SQL databases fills the gap

Brings OLTP support into the Fabric ecosystem—ideal for transactional and operational workloads.

Preserves your SQL logic (stored procedures, triggers, views) with minimal rework

Familiar tools like SQL Server Management Studio (SSMS), Azure Data Studio, and T-SQL are fully supported.

Batabase Projects & Git integration enabling full version control and CI/CD pipelines.

SQL database in Microsoft Fabric is still in **Preview**

Fabric as a unified data platform

With Lakehouse, Warehouse, and SQL Database, Microsoft Fabric now supports the full spectrum: analytics, engineering, and operations.

SQL professionals can now work across all layers of the data stack— without rewriting logic or changing tools.

(•) (•) (*) (*) (*) (•) (•)



Fabric SQL databases fills the gap

Brings OLTP support into the Fabric ecosystem—ideal for transactional and operational workloads.

💩 Preserves your SQL logic (stored procedures, triggers, views) with minimal rework

Familiar tools like SQL Server Management Studio (SSMS), Azure Data Studio, and T-SQL are fully supported.

Lotabase Projects & Git integration enabling full version control and CI/CD pipelines.

1 Tooling support (e.g., Visual Studio, Azure Data Studio) has not fully caught up with the latest updates in Fabric

All operations supported

 \square

Feature	Lakehouse Endpoint	Warehouse	SQL Database
CREATE TABLE	×		
ALTER/DROP TABLE	×		
INSERT/UPDATE/DELETE	×		
CTAS (CREATE AS SELECT)	×		
Identity Columns	×	🗙 (Workaround)	
Temp Tables / CTEs	Views only		
Data Types	Mirrors Delta Lake	Subset T-SQL	Full T-SQL
Read-only vs Read-write	Read-only	Read-write	Read-write
Endpoint availability	Auto for each LH	1 warehouse/item	Auto + analytics endpoint





Demo 02 - Automate

Steps:

- Using FabricCatalyst ⁽¹⁾ I will create a Fabric Workspace with a Lakehouse, Warehouse & SQL Database
- 2. Generate a Backup Package (BACPAC)
- 3. Create a SQL project using Visual Studio and import the schema from existing DB (DACPAC)

Migrate or Mirror: Two Paths to Bring Your Data into Fabric

Migrate

- Small to Medium Databases with Minor Disruption Acceptable
- Legacy Databases Needing Modernization
- Low-Traffic Applications
- Data Warehousing or Analytics Initiatives

Mirror

- High-Availability Databases
- Real-Time Analytics Workloads
- Incremental Adoption of Fabric
- Multi-Team Data Sharing
- Dynamic Schemas with Frequent Changes

Two Key Options for Migration

Fabric adapts well-**Use Case Tool/Format** Notes established SQL Server techniques Schema-only .dacpac + Ideal for version migrations SqlPackage.exe controlled Even as Microsoft Fabric deployments introduces a modern, cloud-native experience, Schema + data Suitable for full .bacpac + but they are adapting well-SqlPackage.exe migration database established SQL Server migrations techniques

(•) (•) (*) (*) (*) (*) (*)

1 Tooling support (e.g., Visual Studio, Azure Data Studio) has not fully caught up with the latest updates in Fabric



Demo 03 - Migrate

Steps:

- Using a DevOps pipeline, I'll execute SqlPackage ⁽¹⁾ to deploy a BACPAC into Fabric's SQL Database
- Using a DevOps pipeline, I'll execute SqlPackage ⁽¹⁾ to deploy a DACPAC into Fabric's SQL Database
- 3. Get the workspace Git-enabled



Demo 04 - Mirror

Steps:

- 1. Create a Connection
- 2. Verify SKU is at least 100 DTU's
- 3. Create a mirrored Database



TechTacoFriday Technology spiced with a taco flavor mix

Thank you for your attention, see you around © All the best



24