



SQL Saturday Toronto 2025



Our Sponsors



Toronto Data Professionals Community (TDPC)



Community Support

[Toronto Data Professionals Community \(TDPC\)](#), one of the largest data professional's community in Toronto, host monthly event which offers interactive learning built by community and guided by trusted data experts.

TDPC Event Partners



Toronto Data Professionals Community (TDPC)




SQL Saturday (#1131)



SQL in Microsoft Fabric: A Unified Approach to Data Engineering, Modelling, and Intelligence

Alexander Turlov
Principal consultant



Toronto Data Professionals Community (TDPC)



SQL Saturday (#1131)

Alexander Turlov

Principal Consultant | Alta Computing Inc.



Helping organizations modernize apps, cloud and data platforms with
Microsoft Azure & Fabric

Industry certified



Trusted by Partners



Trusted by Customers



Toronto Data Professionals Community (TDPC)



SQL Saturday (#1131)

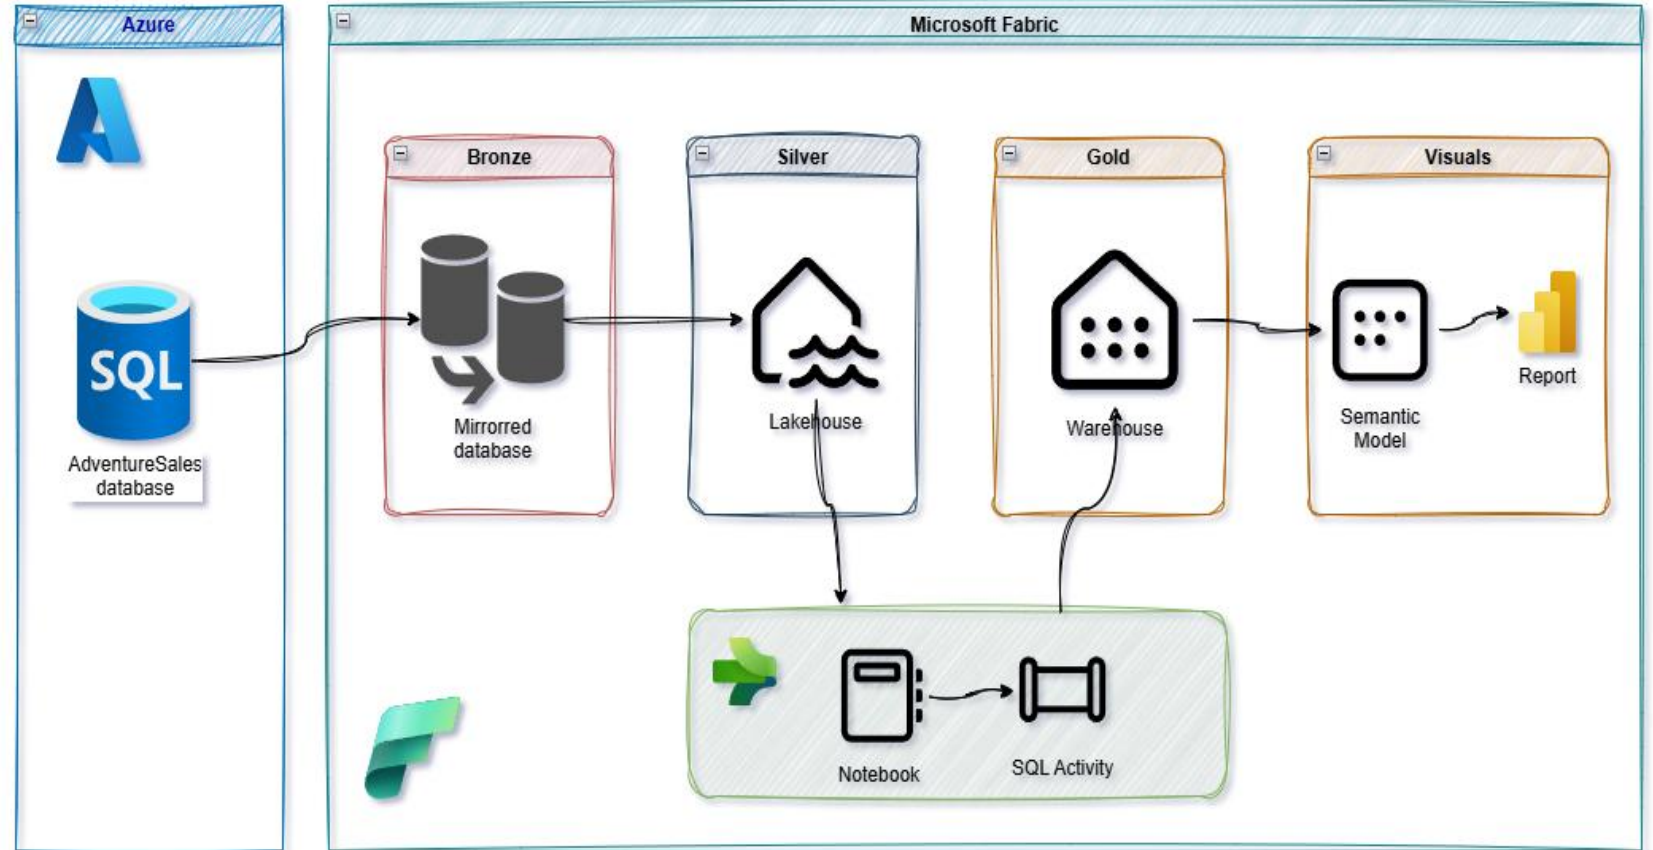
Why SQL in Microsoft Fabric

- Unified experience across **Lakehouse, Warehouse, Pipelines**
- SQL powers **data engineering, modeling, BI**
- Familiar syntax → **fast adoption**
- Today: **Retail Sales Analytics demo-first**



Solution Architecture

Bronze: Raw data
Silver: Curated data
Gold: Analytics-ready data
Power BI: Direct Lake mode for instant visuals



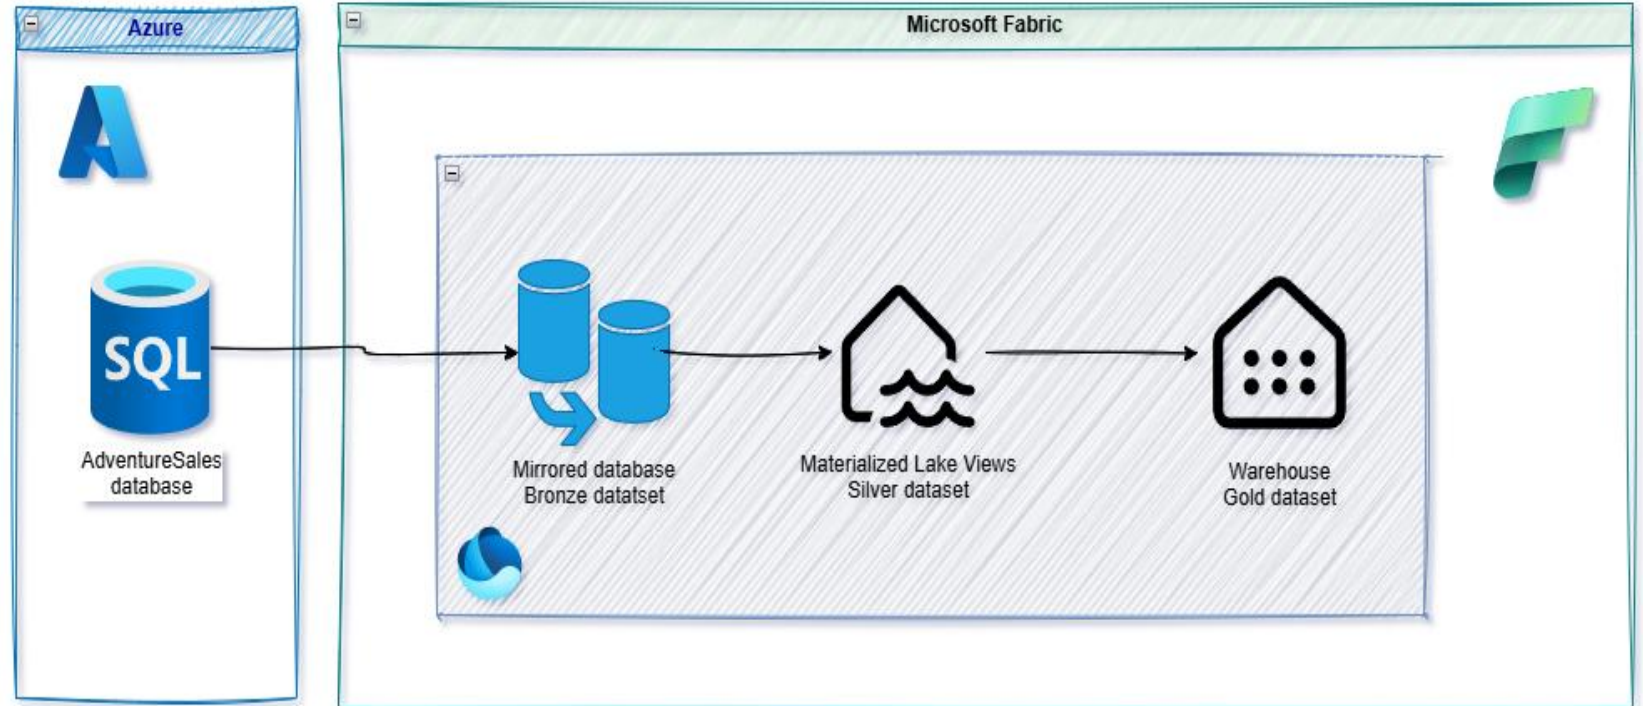
Data Flow

Bronze: Automatic replication

Silver: Refresh MLVs from a Notebook

Gold: Incremental load by a pipeline

Power BI: automatic refresh via Direct Lake mode



How the Demo Works

- Examine Fabric items starting with Power BI report, Semantic model, Warehouse, Lakehouse and Mirrored database.
- Simulate data update
- Trigger the data refresh in Fabric and observe flow of changes
- Confirm the data refresh



How Incremental Load Works

1. Source tables need **ModifiedDate** column for change capture.
2. Maintain a metadata table to track refresh **watermarks**.
3. When MLVs refresh they capture **ModifiedDate** values from the source tables.
4. When Warehouse loaders execute, they filter rows by > **ModifiedDate**.
5. Update **watermark** value in the control table after load.

```
-- 1) Read watermark
DECLARE @w DATETIME2 =
    (SELECT last_watermark FROM etl.watermark WHERE src = 'sales');

-- 2) Upsert only changed rows
MERGE gold.fact_sales AS t
USING (
    SELECT SalesOrderID, SalesOrderDetailID, TotalDue, ModifiedDate
    FROM silver.mlv_sales
    WHERE ModifiedDate > @w
) AS s
ON (t.SalesOrderID = s.SalesOrderID AND t.SalesOrderDetailID = s.SalesOrderDetailID)
WHEN MATCHED THEN
    UPDATE SET t.TotalDue = s.TotalDue, t.ModifiedDate = s.ModifiedDate
WHEN NOT MATCHED THEN
    INSERT (SalesOrderID, SalesOrderDetailID, TotalDue, ModifiedDate)
    VALUES (s.SalesOrderID, s.SalesOrderDetailID, s.TotalDue, s.ModifiedDate);

-- 3) Advance watermark
UPDATE etl.watermark
SET last_watermark = (SELECT MAX(ModifiedDate) FROM silver.mlv_sales)
WHERE src = 'sales';
```



Thank You!



Alexander Turlov

Principal Consulting
Alta Computing
aturlov@altacomputing.com

Post-event feedback survey



Toronto Data Professionals Community (TDPC)



SQL Saturday (#1131)