#FABCONSQLCON2026

FABCON
Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON
Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026

The Ultimate “NO ETL” Showdown

Shortcuts vs. Mirroring vs. Copy Job

Nikola Ilic | Data Mozart | Microsoft Data Platform MVP

Your “No ETL” Showdown Host

Principal Architect

iLink Digital

Nikola Ilic

- *I'm making music from the data*!
- Power BI and Fabric addict, blogger, speaker…
- Father of 2, Barca & Leo Messi fan…

data-mozart.com

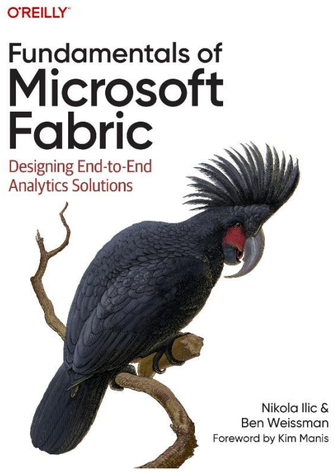

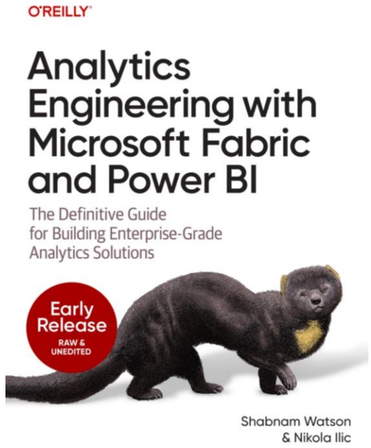

PLURALSIGHT

O'REILLY®

You are paying an ETL tax…Let’s end that today!

Every hour spent wiring pipelines, debugging flows, and managing schedules is an hour NOT spent on actual insights!

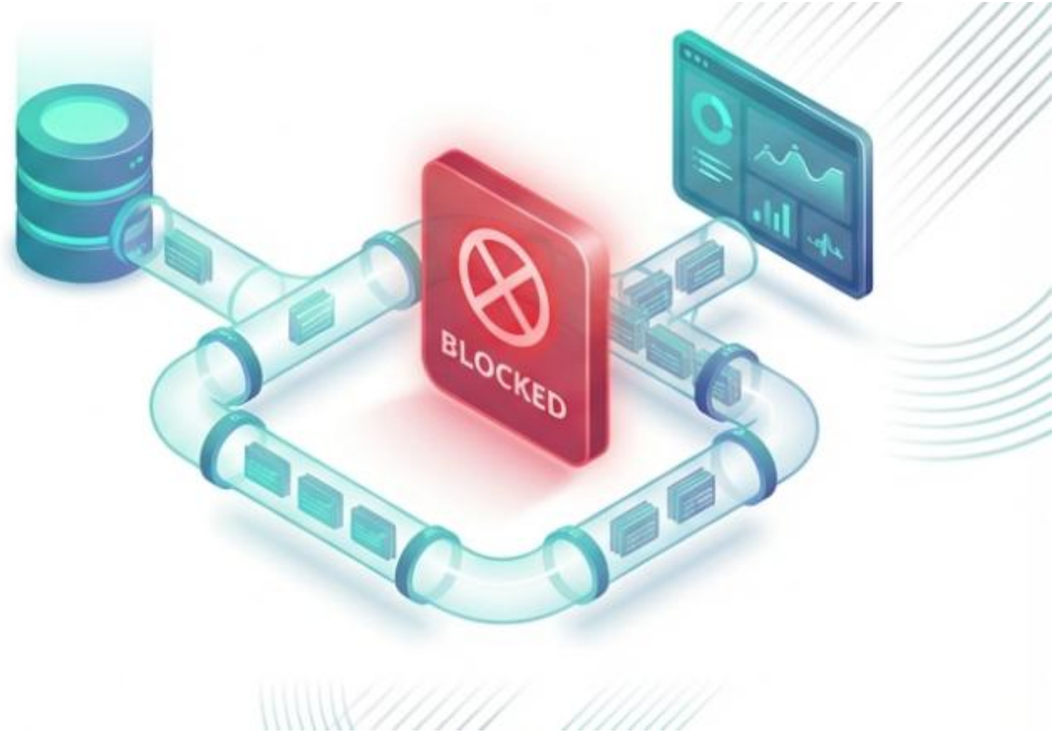

The “NO ETL” promise

Fabric provides three distinct ways to access external data without building pipelines.

Three approaches -> One goal

Stop buying the DVD – just press PLAY!

What Does “No ETL” Really Mean?

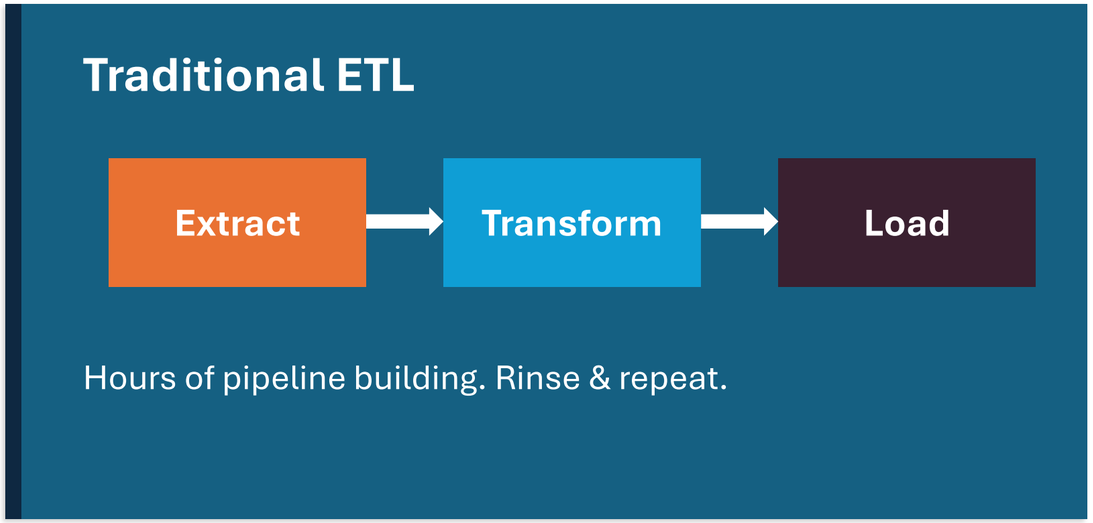

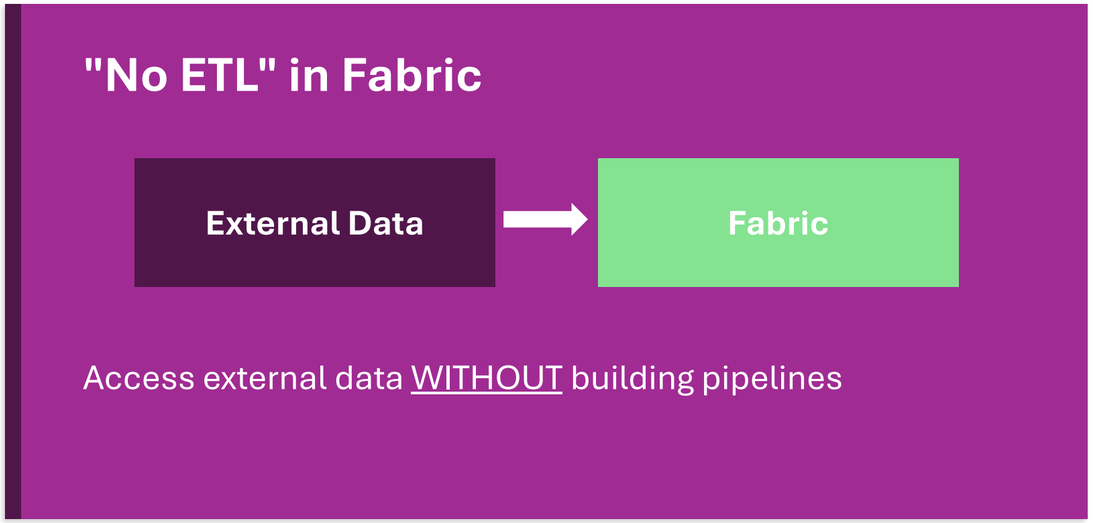

Meet the Contestants

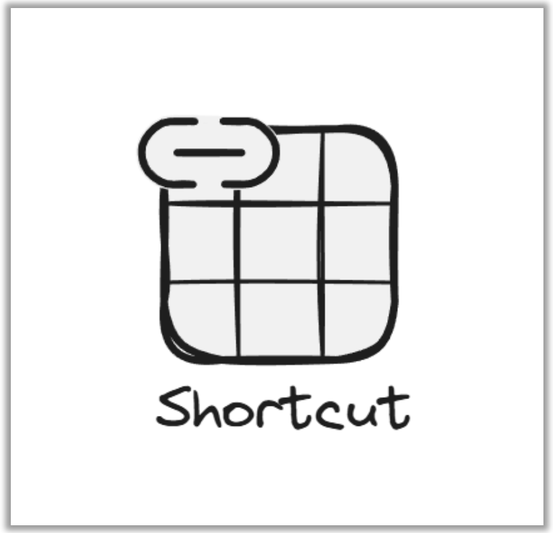

Three fighters. One shared goal. Only YOUR use case decides the winner.

Contestant 1: Shortcuts

Bookmark

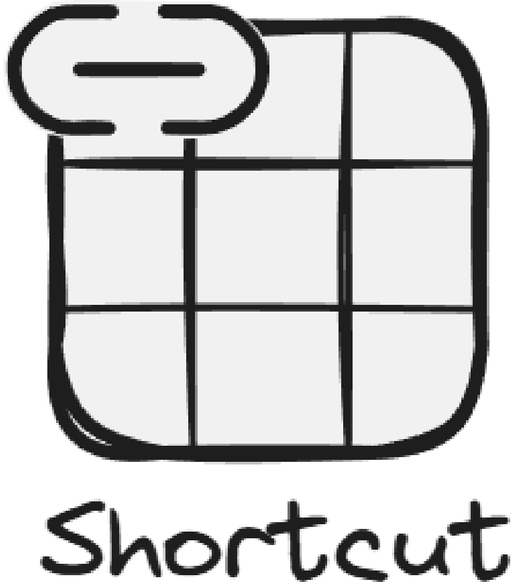

- ✓ Zero-copy symbolic links. Data never moves
- ✓ Like streaming a movie: zero download, just press play
- ✓ Sources
  - ✓ ADLS Gen2
  - ✓ Amazon S3
  - ✓ GCS
  - ✓ Dataverse
  - ✓ SharePoint
  - ✓ OneDrive
- ✓ Cost: Egress fees only

Contestant 1: Shortcuts

Bookmark

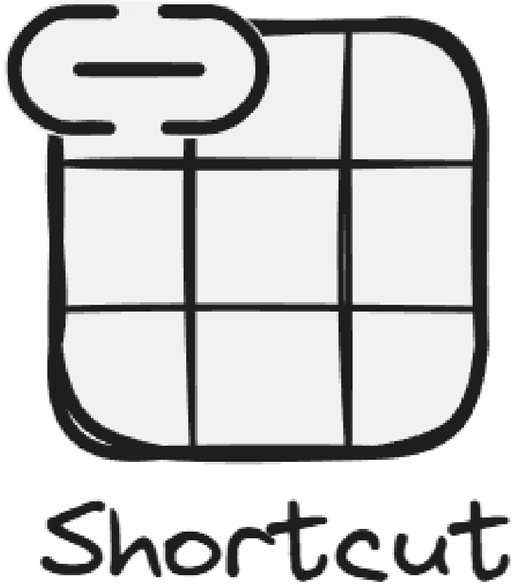

- ✓ Zero-copy symbolic links. Data never moves
- ✓ Like streaming a movie: zero download, just press play
- ✓ Sources
  - ✓ ADLS Gen2
  - ✓ Amazon S3
  - ✓ GCS
  - ✓ Dataverse
  - ✓ SharePoint
  - ✓ OneDrive
- ✓ Cost: Egress fees only

How Shortcuts work?

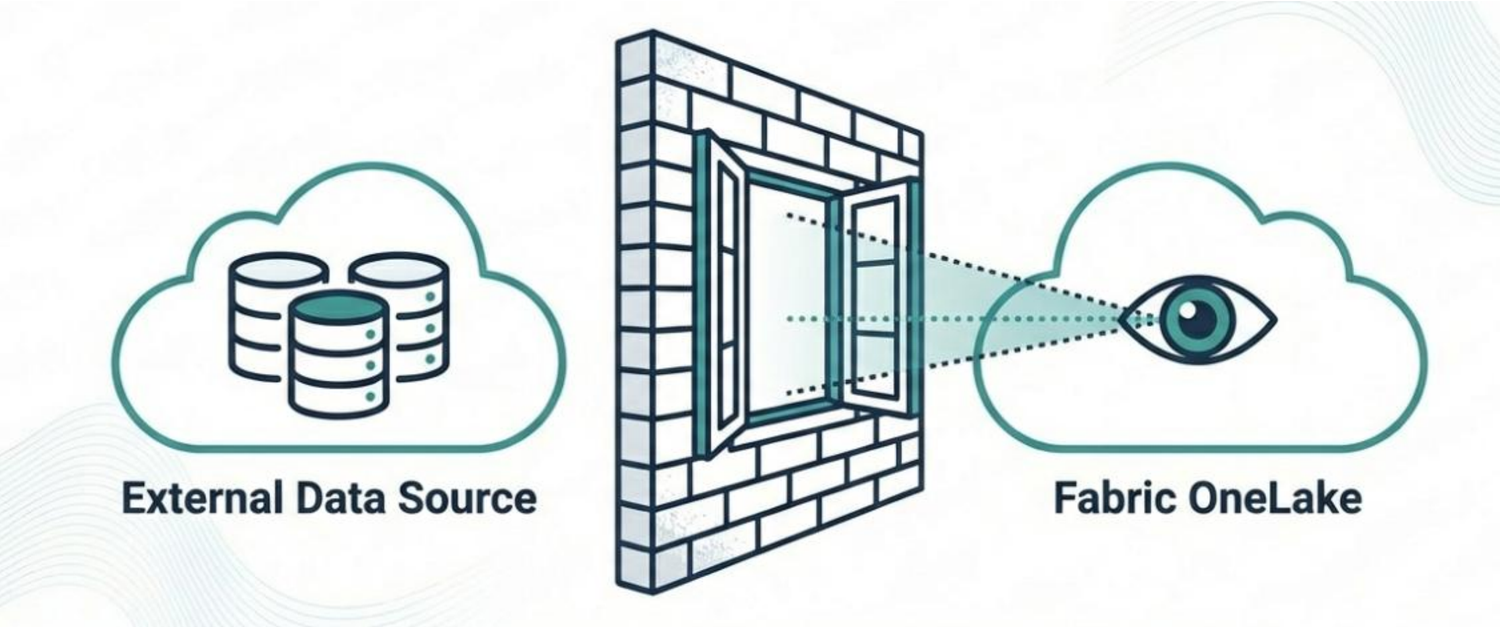

The data never moves. Fabric looks through a window to see your data on the other side

External Shortcuts in a nutshell

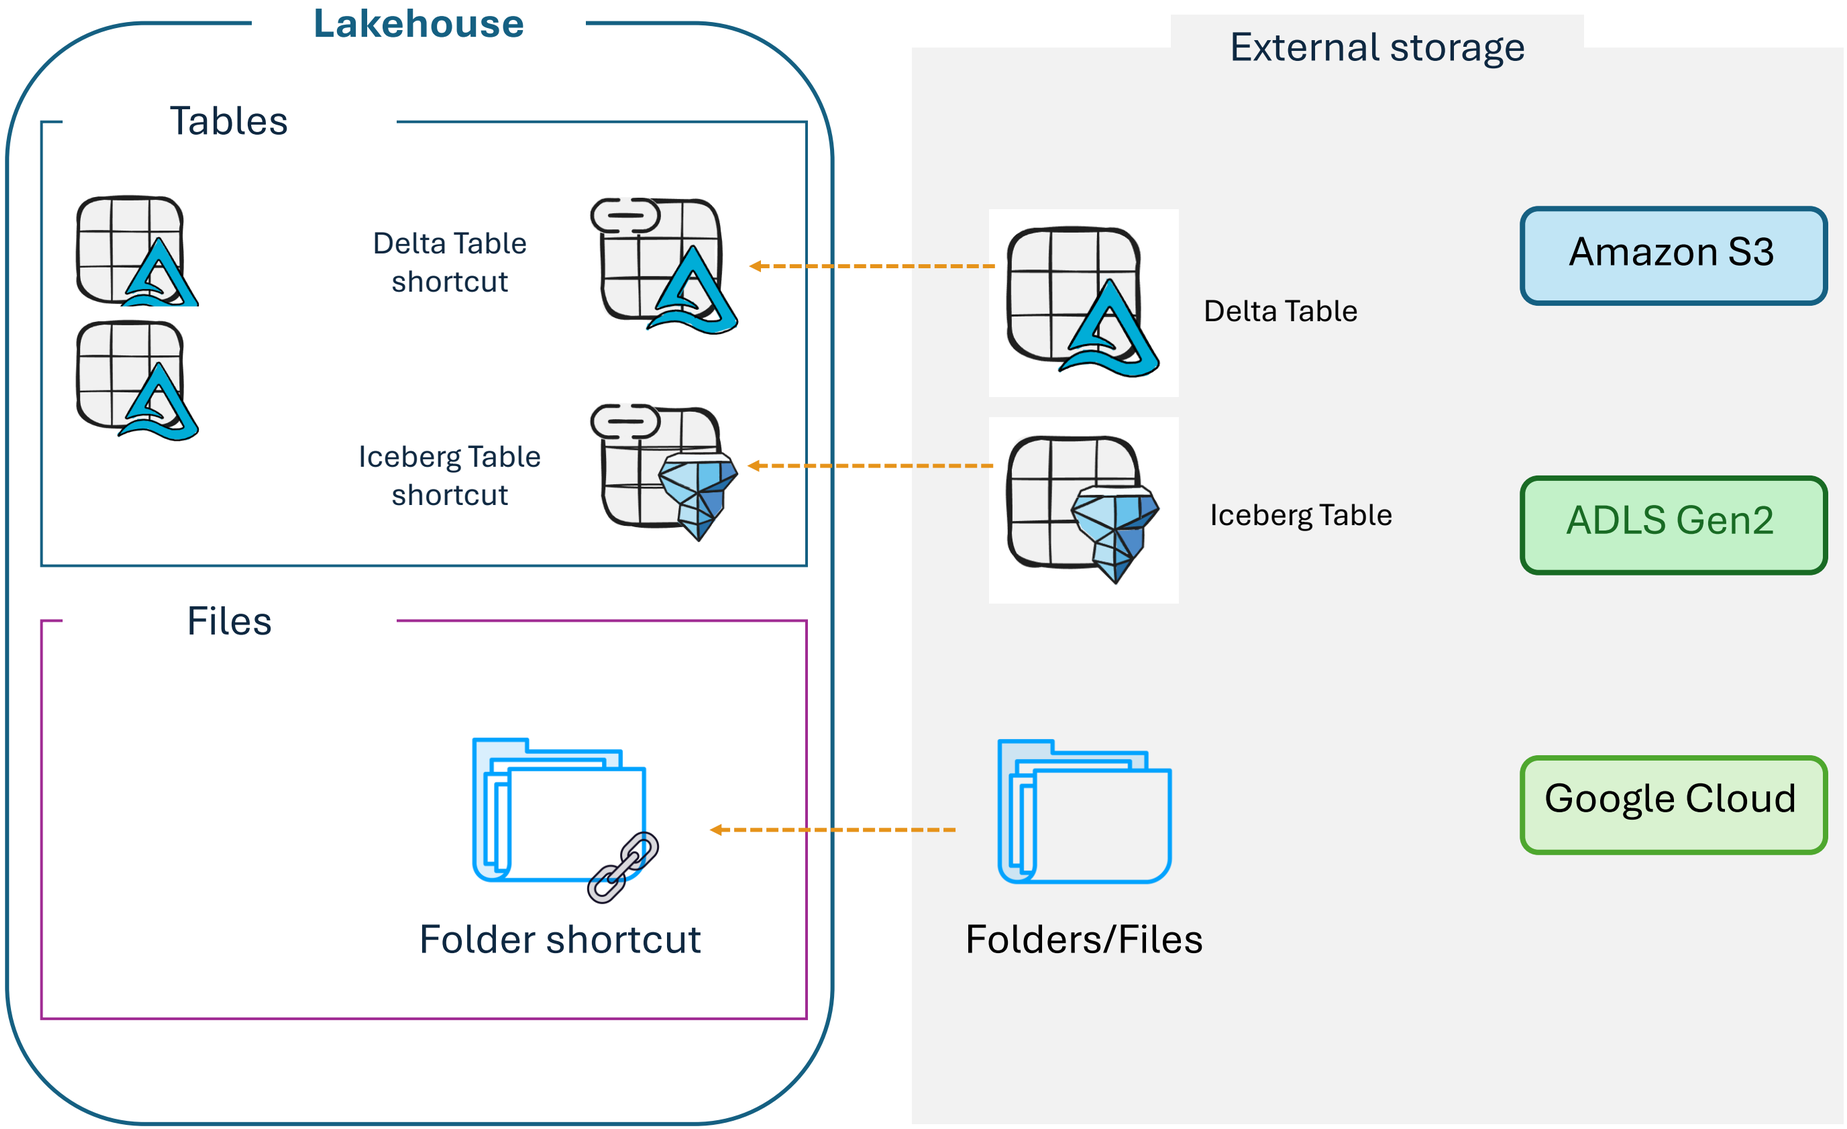

Internal Shortcuts

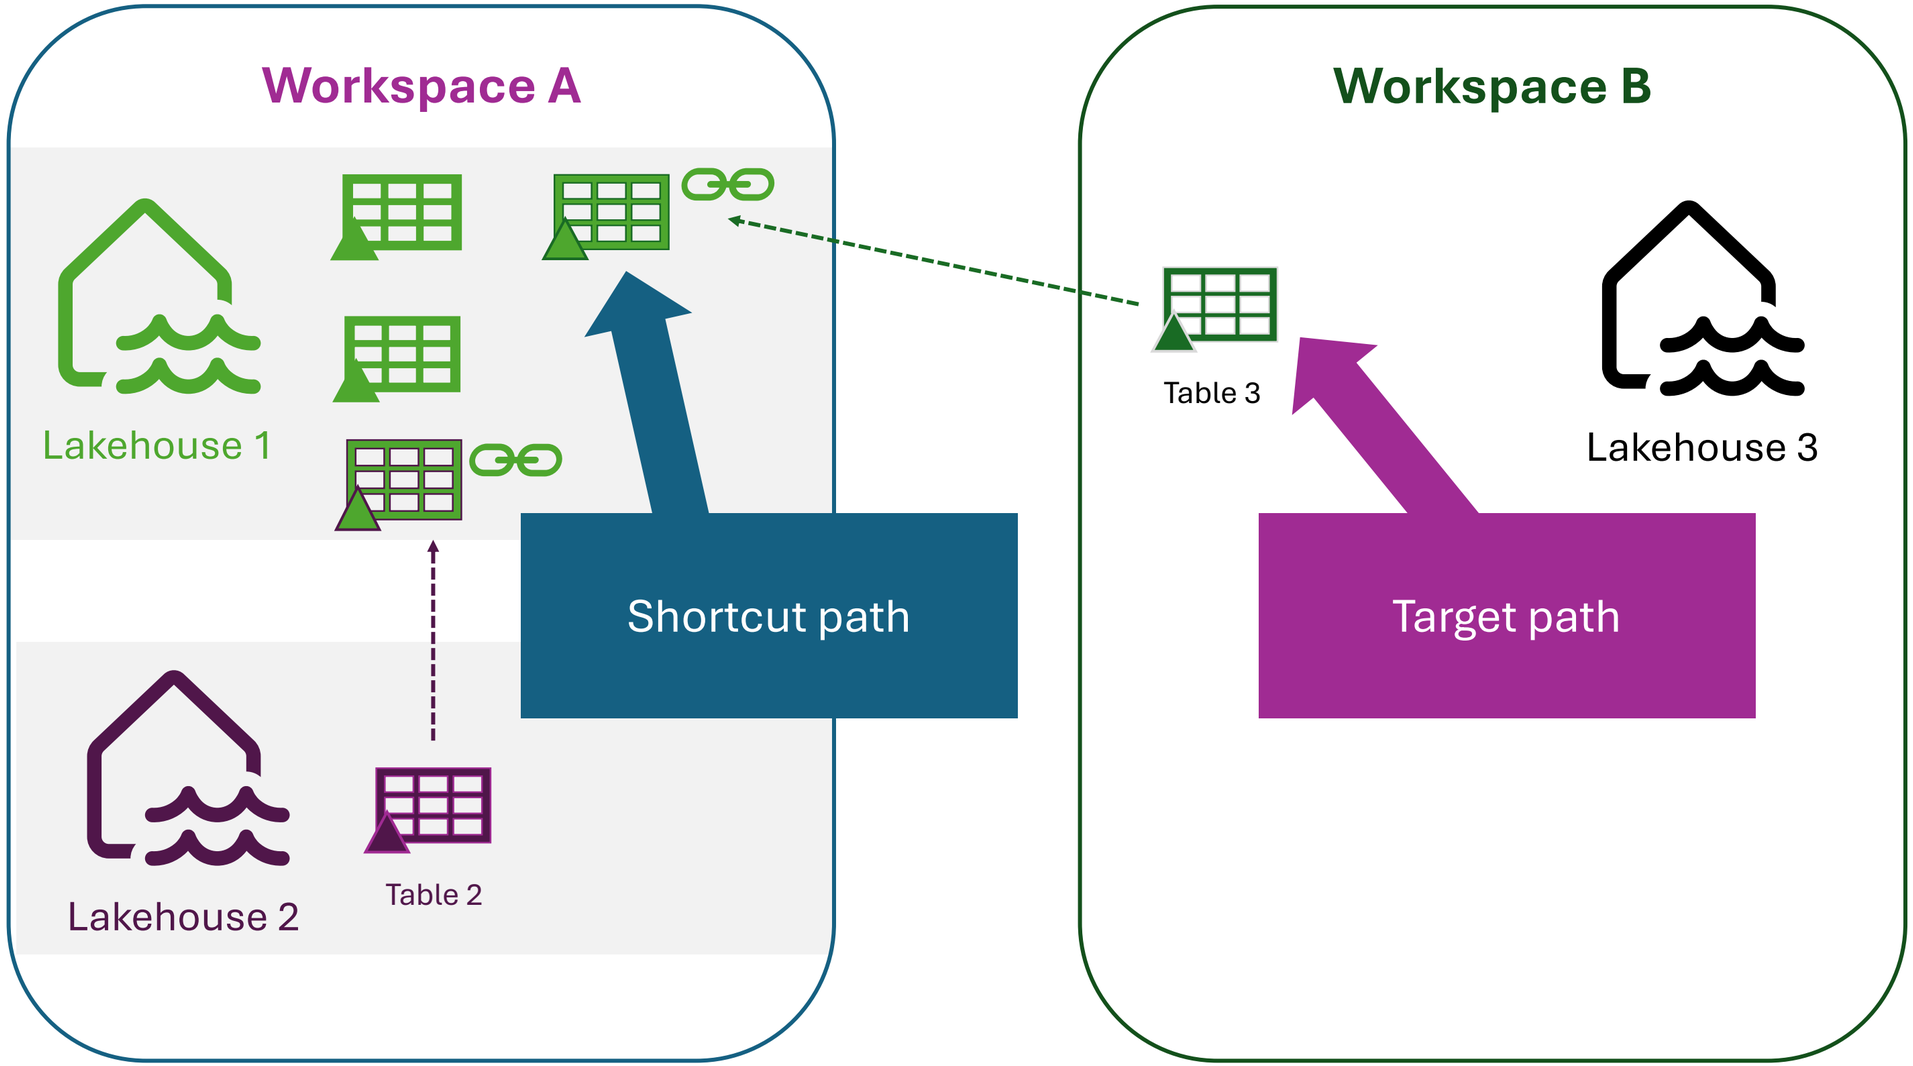

Shortcuts: The Good, the Bad, and the Gotchas

THE GOOD

- Zero data movement
- Instant setup in seconds
- No storage cost in Fabric
- Works with Direct Lake
- Up to 100K per item

THE BAD

- Performance depends on external source
- Cannot shortcut to databases
- No transformations possible
- Egress fees from cloud providers

GOTCHAS

- Files >1GB not cached (cross-cloud)
- No non-Latin chars in names
- Max 5 shortcut-to-shortcut depth
- Caching only for S3/GCS

Demo

Shortcuts in action

Creating a shortcut to ADLS Gen2 and querying from the lakehouse

Contestant 2: Mirroring

The live photocopier

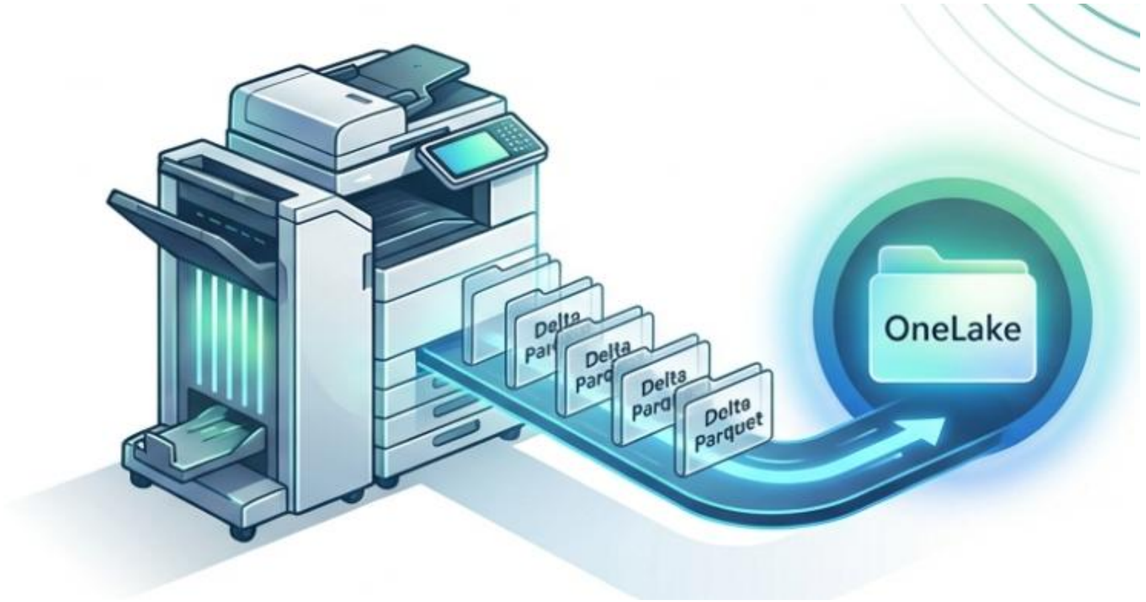

- ✓ Continuous – Change Data Capture (CDC) only copies what changed
- ✓ Three flavors – Database, Metadata (Databricks), Open
- ✓ Cost – Replication compute in Fabric is FREE (and some of the storage☺)

Mirroring: Three Unique Flavors

Database Mirroring

Full replication of tables into OneLake as Delta

Azure SQL DB, SQL MI, SQL Server 2016-2025, Snowflake, Cosmos DB, PostgreSQL

Metadata Mirroring

Syncs catalog structure ONLY, NOT actual data

Azure Databricks Unity Catalog (hybrid: metadata + shortcuts)

Open Mirroring

Any app writes CDC data into mirrored database

Oracle GoldenGate, Qlik, Striim, SAP, CData + 10 more partners

The Magic of Change Data Capture (CDC)

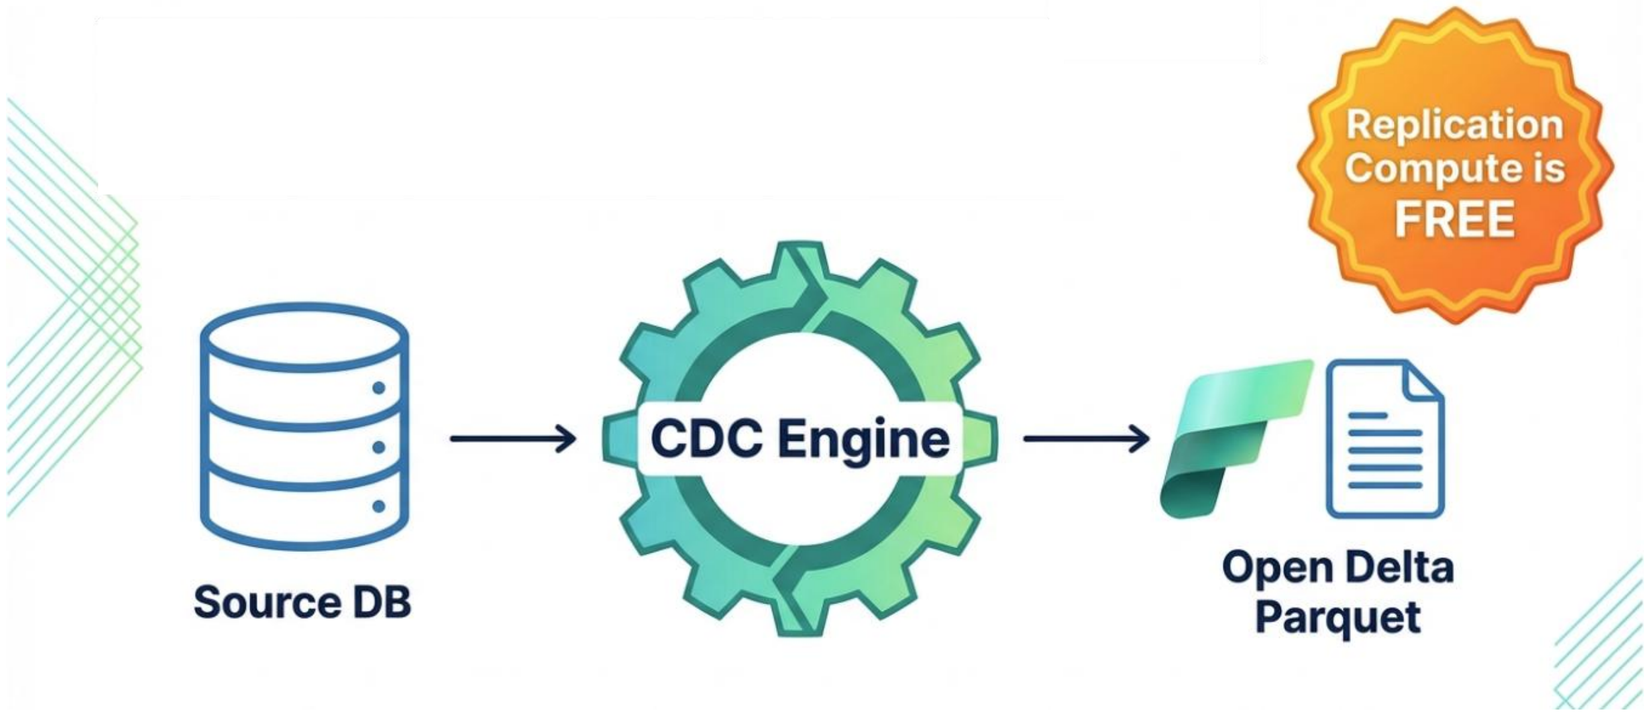

Mirroring doesn’t copy everything every time.

It only captures what changed, converting it instantly into open formats

Mirroring: The Good, the Bad, and the Gotchas

THE GOOD

- Near real-time sync
- FREE replication compute
- Delta format -> Direct Lake
- Predictable local performance
- Auto converts to open format

THE BAD

- Destination is read-only
- Limited to 1,000 tables
- No transformations
- Can't shortcut to databases

GOTCHAS

- No custom scheduling
- Always-on replication
- Stop + restart = full reseed!
- CDC conflict with SQL Server

Demo

Mirroring in action

Setting up database mirroring

Contestant 3: Copy Job

The moving company

- ✓ You dictate what to move, when, and where
- ✓ Scheduled or on-demand physical data copies
- ✓ Over 80+ supported connectors

Complete Control

- ✓ Granular table and column mapping -> rename on the fly
- ✓ Native incremental loads using watermark columns and upserts

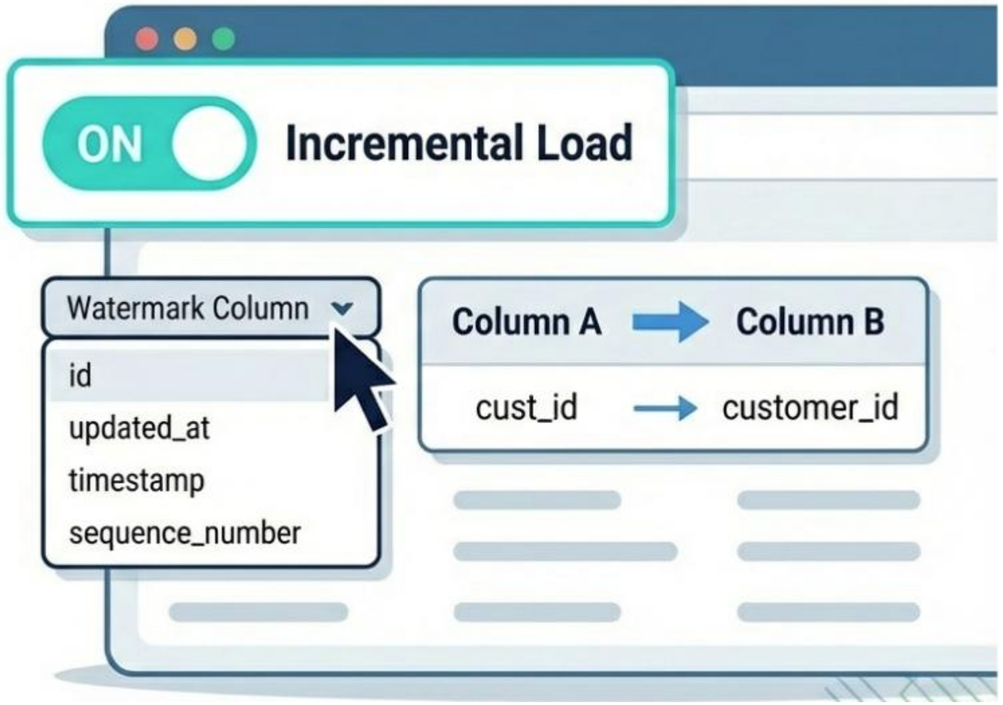

Demo

Copy Job in action

Creating a Copy Job with incremental load

It’s show time…

Round #1: Architecture and movement

Shortcuts: STREAMING

The movie stays on their servers.
Zero data movement

Mirroring: DVR

Auto-records new episodes.
Continuous CDC into OneLake

Copy Job: DOWNLOADS

You pick what and when.
Scheduled batch movement

Round #2: Performance and Cost

	Shortcut	Mirroring	Copy Job
Performance	**Variable** Depends on external source speed	**Fast** Local data in OneLake Direct Lake optimized	**Predictable** Local after copy Schedule-dependent
Cost	**Egress only** No Fabric storage Cloud provider fees	**FREE Compute!** Storage based on capacity SKU	**CU Based** Pay for compute during copy runs

For raw query speed: Local data always wins (Mirroring & Copy Job)

The Comparison Matrix

>>> Screenshot this!

Feature	Shortcuts	Mirroring	Copy Job
Data movement	None (pointer)	Continuous CDC	Scheduled
Latency	Real-time*	Near real-time	Schedule-based
Sources	Storage only	Databases	80+ connectors
Destination	OneLake (virtual)	OneLake (read-only)	Multiple targets
Cost	Egress fees only	FREE compute!	CU consumption
Scheduling	N/A	Always-on	Custom
Flexibility	Low	Low	High
Direct Lake	Yes (external perf.)	Yes (optimized!)	Yes (after copy)

* Shortcuts real-time depends on external system performance

The Decision Framework

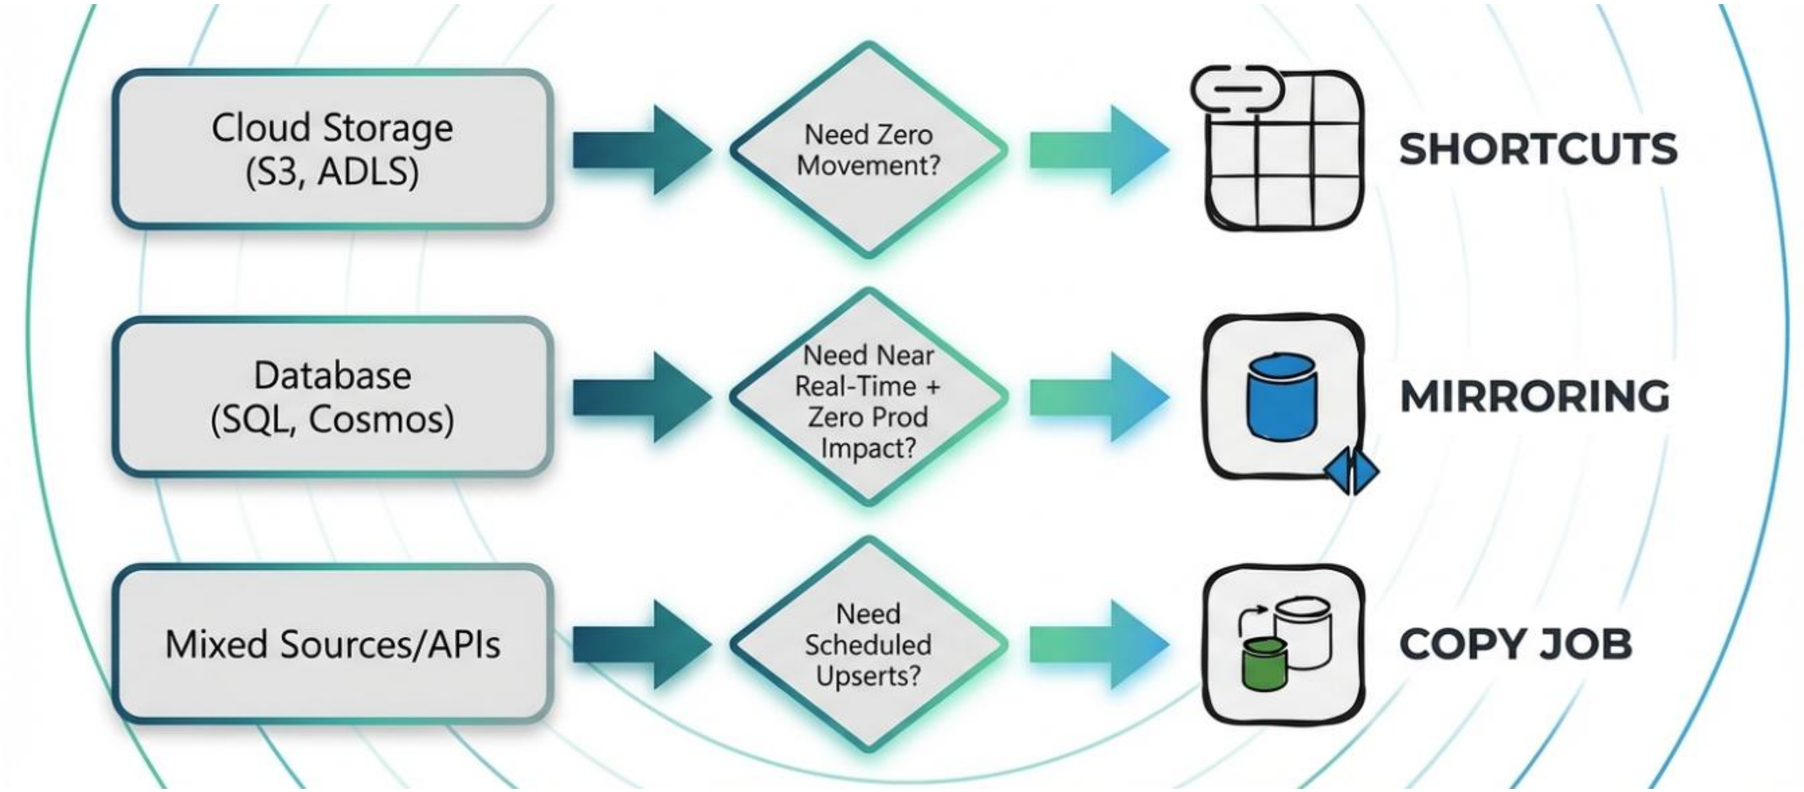

Scenario #1: James, the Finance Manager

James the Finance Manager

Scenario

- Azure SQL DB with policy data
- Needs a real-time executive dashboards
- Can’t slow down production
- Not technical, budget-conscious

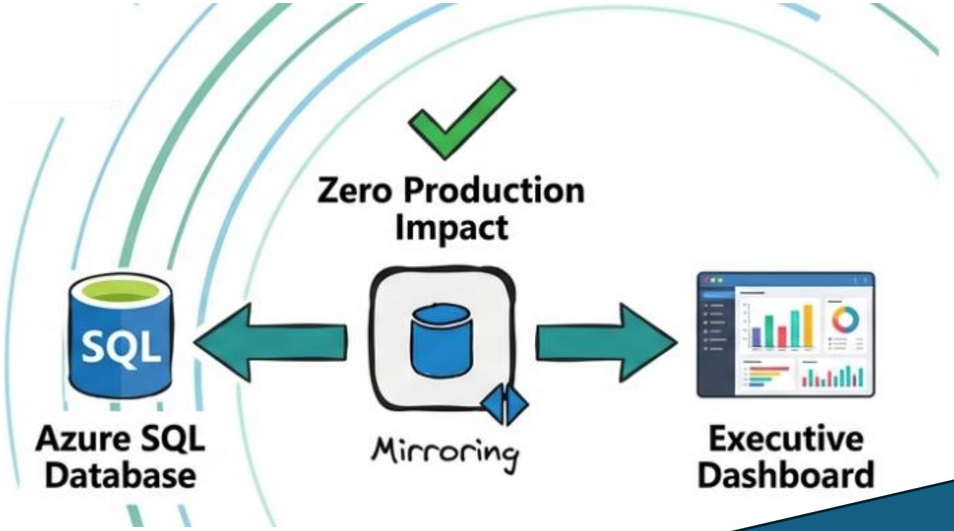

The answer: Mirroring

James is happy. The CFO is happy. Nobody's database is on fire

Scenario #2: Sarah, the Cloud Architect

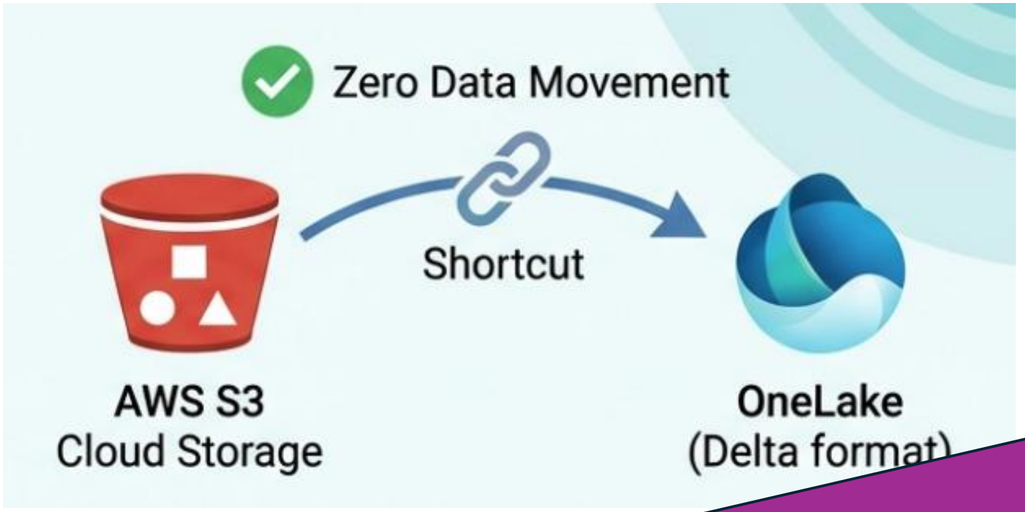

Scenario

- 50 TB data lake in AWS S3
- Multi-cloud strategy
- Can’t afford to copy 50 TB
- Needs Fabric analytics on S3 data

The answer: Shortcuts

Sarah saves on storage and still gets the full Fabric analytics experience

Scenario #3: Lisa, the Business Analyst

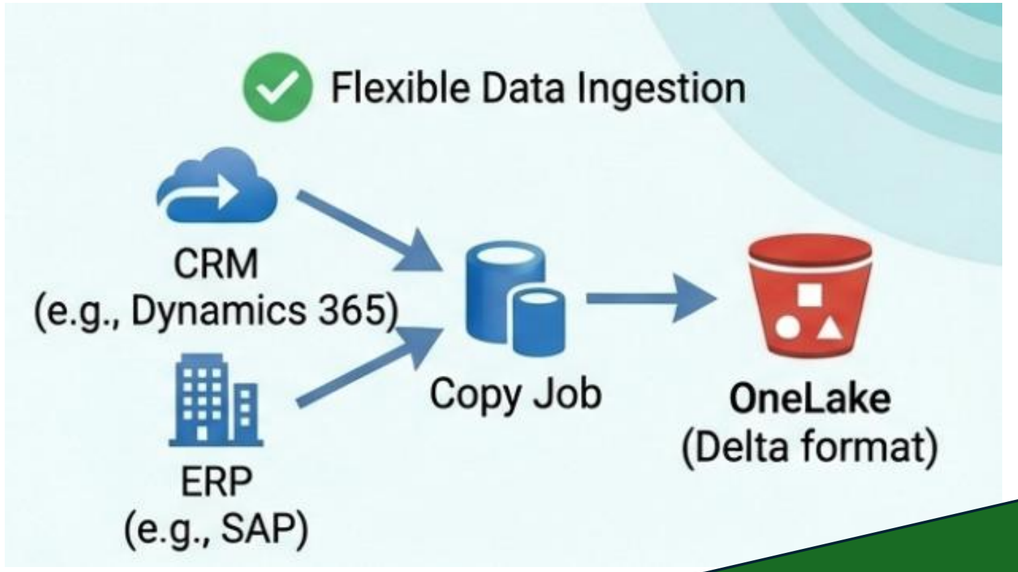

Scenario

- Shipment data from Snowflake
- Custom schedule: every 4 hours
- Needs column mapping and upserts
- Multiple source databases

The answer: Copy Job

Lisa gets supply chain analysis without becoming a data engineer

Plot Twist: You Don’t Have to Choose Just One!

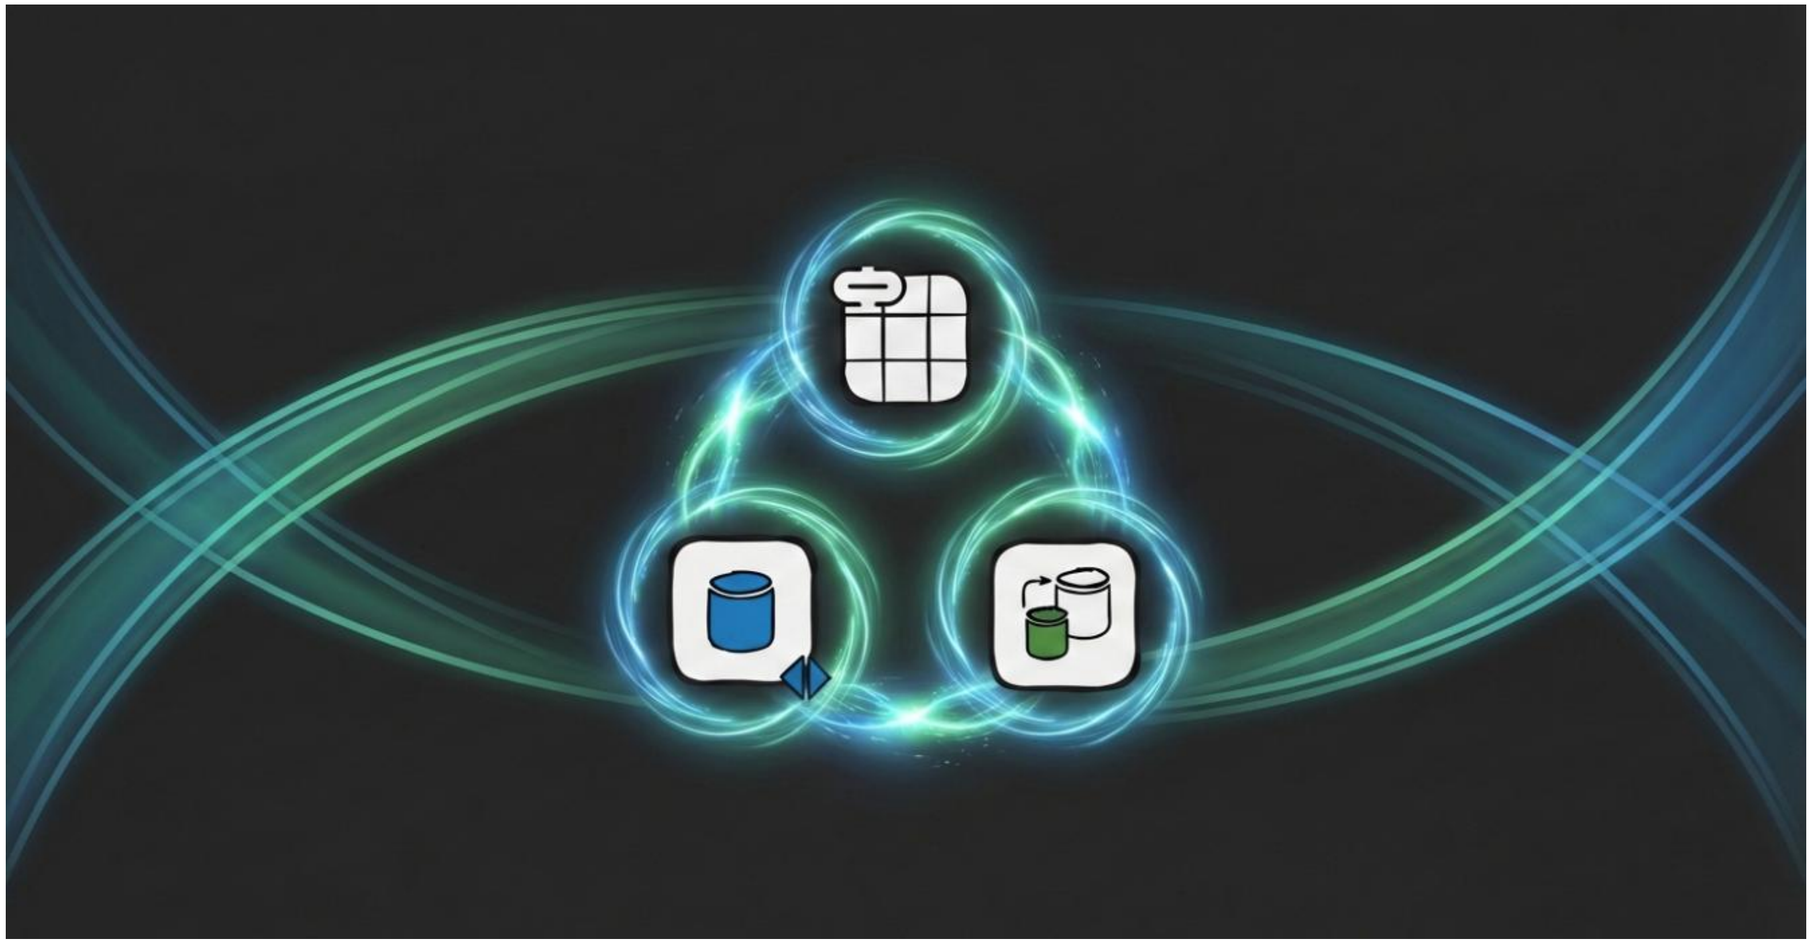

Match the tool to the SOURCE TYPE and your LATENCY needs. There is no best tool - only the RIGHT tool.

The Unified “No ETL” Architecture

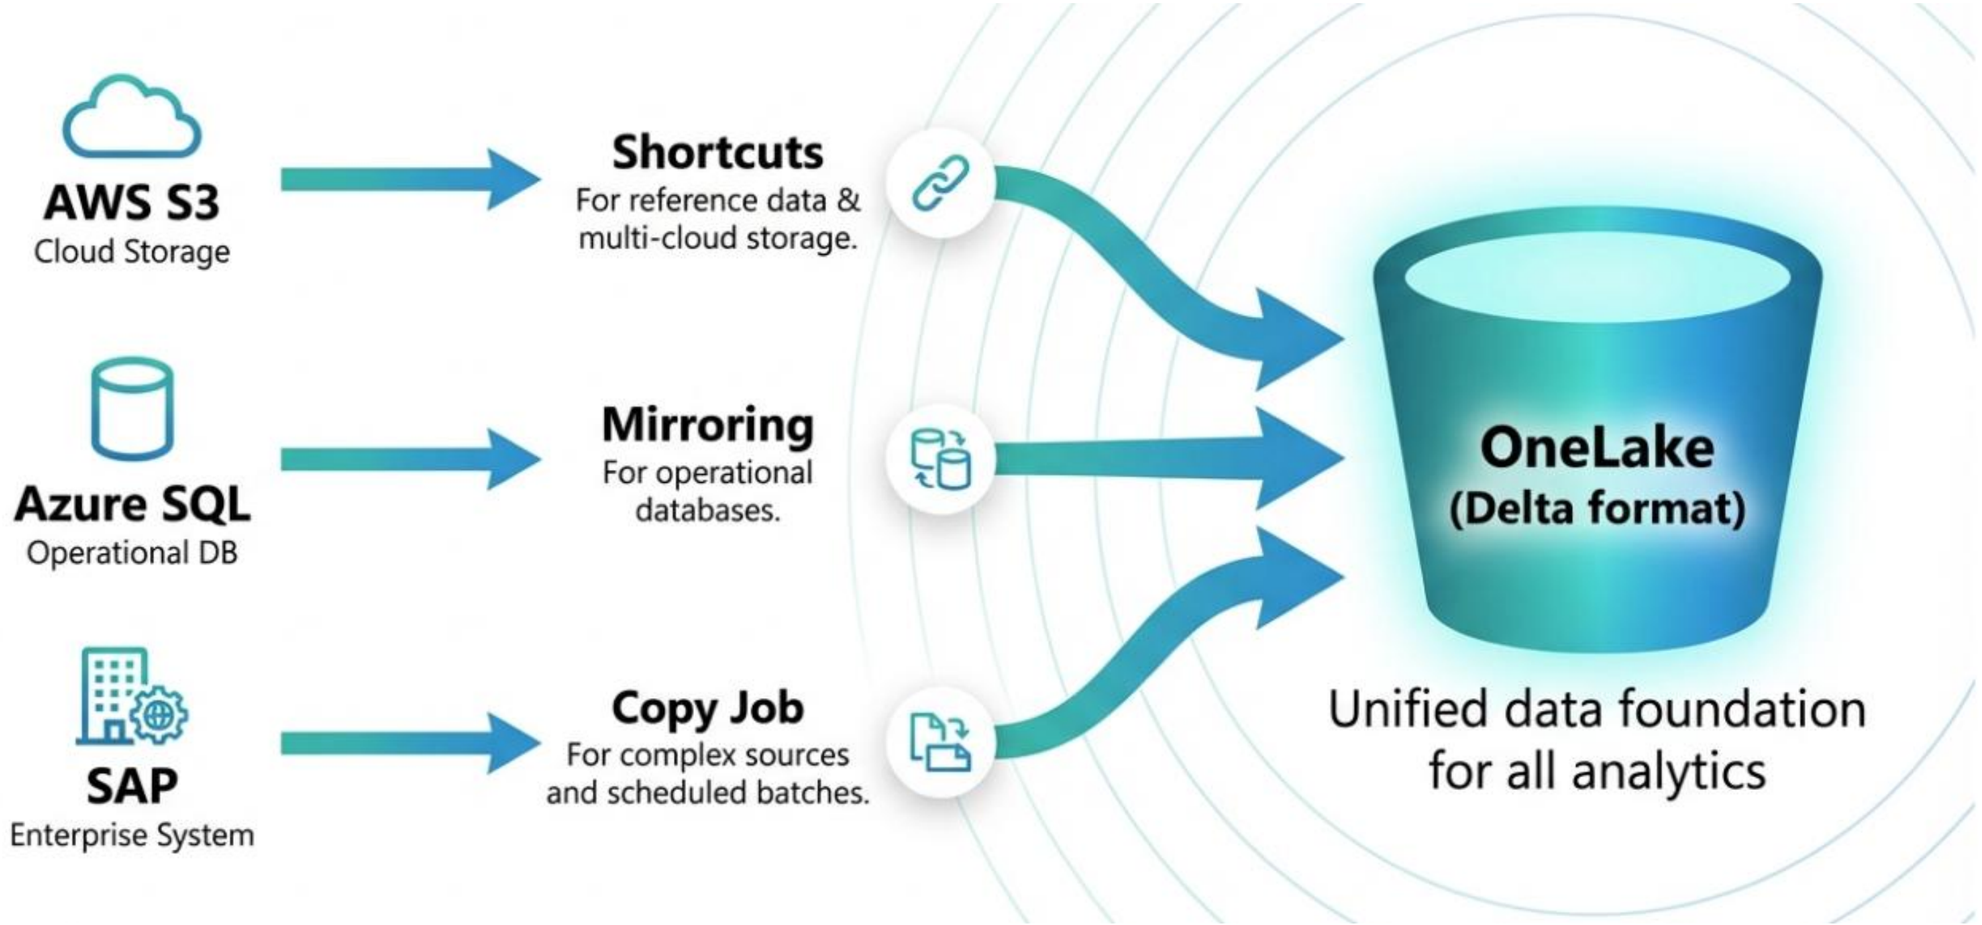

Sound off.
The mic is all yours.
Influence the product roadmap.

Join the Fabric User Panel

Share your feedback directly with our Fabric product group and researchers.

https://aka.ms/JoinFabricUserPanel

Join the SQL User Panel

Influence our SQL roadmap and ensure it meets your real-life needs

https://aka.ms/JoinSQLUserPanel

How was the session?

Complete Session Surveys in Whova for your chance to WIN PRIZES!

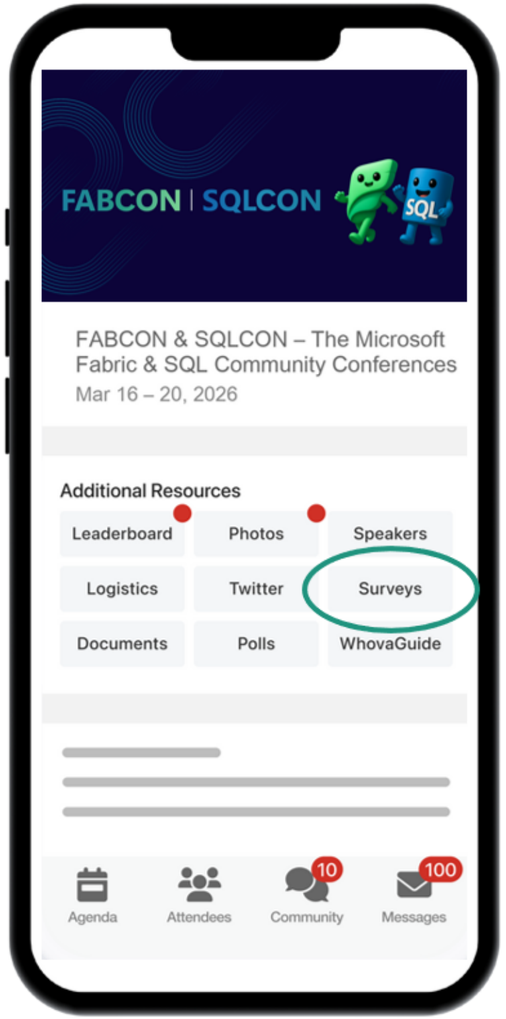

Get Two Fabric Certifications for FREE

Attendees of FABCON can take the Fabric Analytics Engineer or Fabric Data Engineer exam for free. Be part of the 2 fastest growing role-based certifications in Microsoft history.

Request your voucher by March 23, 2026.

https://aka.ms/fabcon/cert100

Thank YOU!