

#FABCONSQLCON2026

FABCON

Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON

Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026



Peachy Migrations: Sweet Moves from Synapse to Fabric

Rob Carrol

Principal Consultant, Coeo

rob.carrol@coeo.com

<https://www.linkedin.com/in/rob-carrol-sql/>



Why Migrate?



Migration tools



Lessons learned



Microsoft Fabric

The unified data platform for AI transformation



Data Factory



Analytics



Databases



Real-Time
Intelligence



IQ



Power BI

Fabric Platform



AI



OneLake



Security & Governance



Microsoft Fabric

The unified data platform for AI transformation

- Data Factory
- Data Engineering
- Data Warehouse
- Data Science
- Databases
- Real-Time Intelligence
- IQ
- Power BI
- Partner Workloads

AI in Microsoft Fabric

OneLake

Security and Governance



Why migrate to **Fabric** Data Warehouse?

Single source of truth, open storage format on OneLake

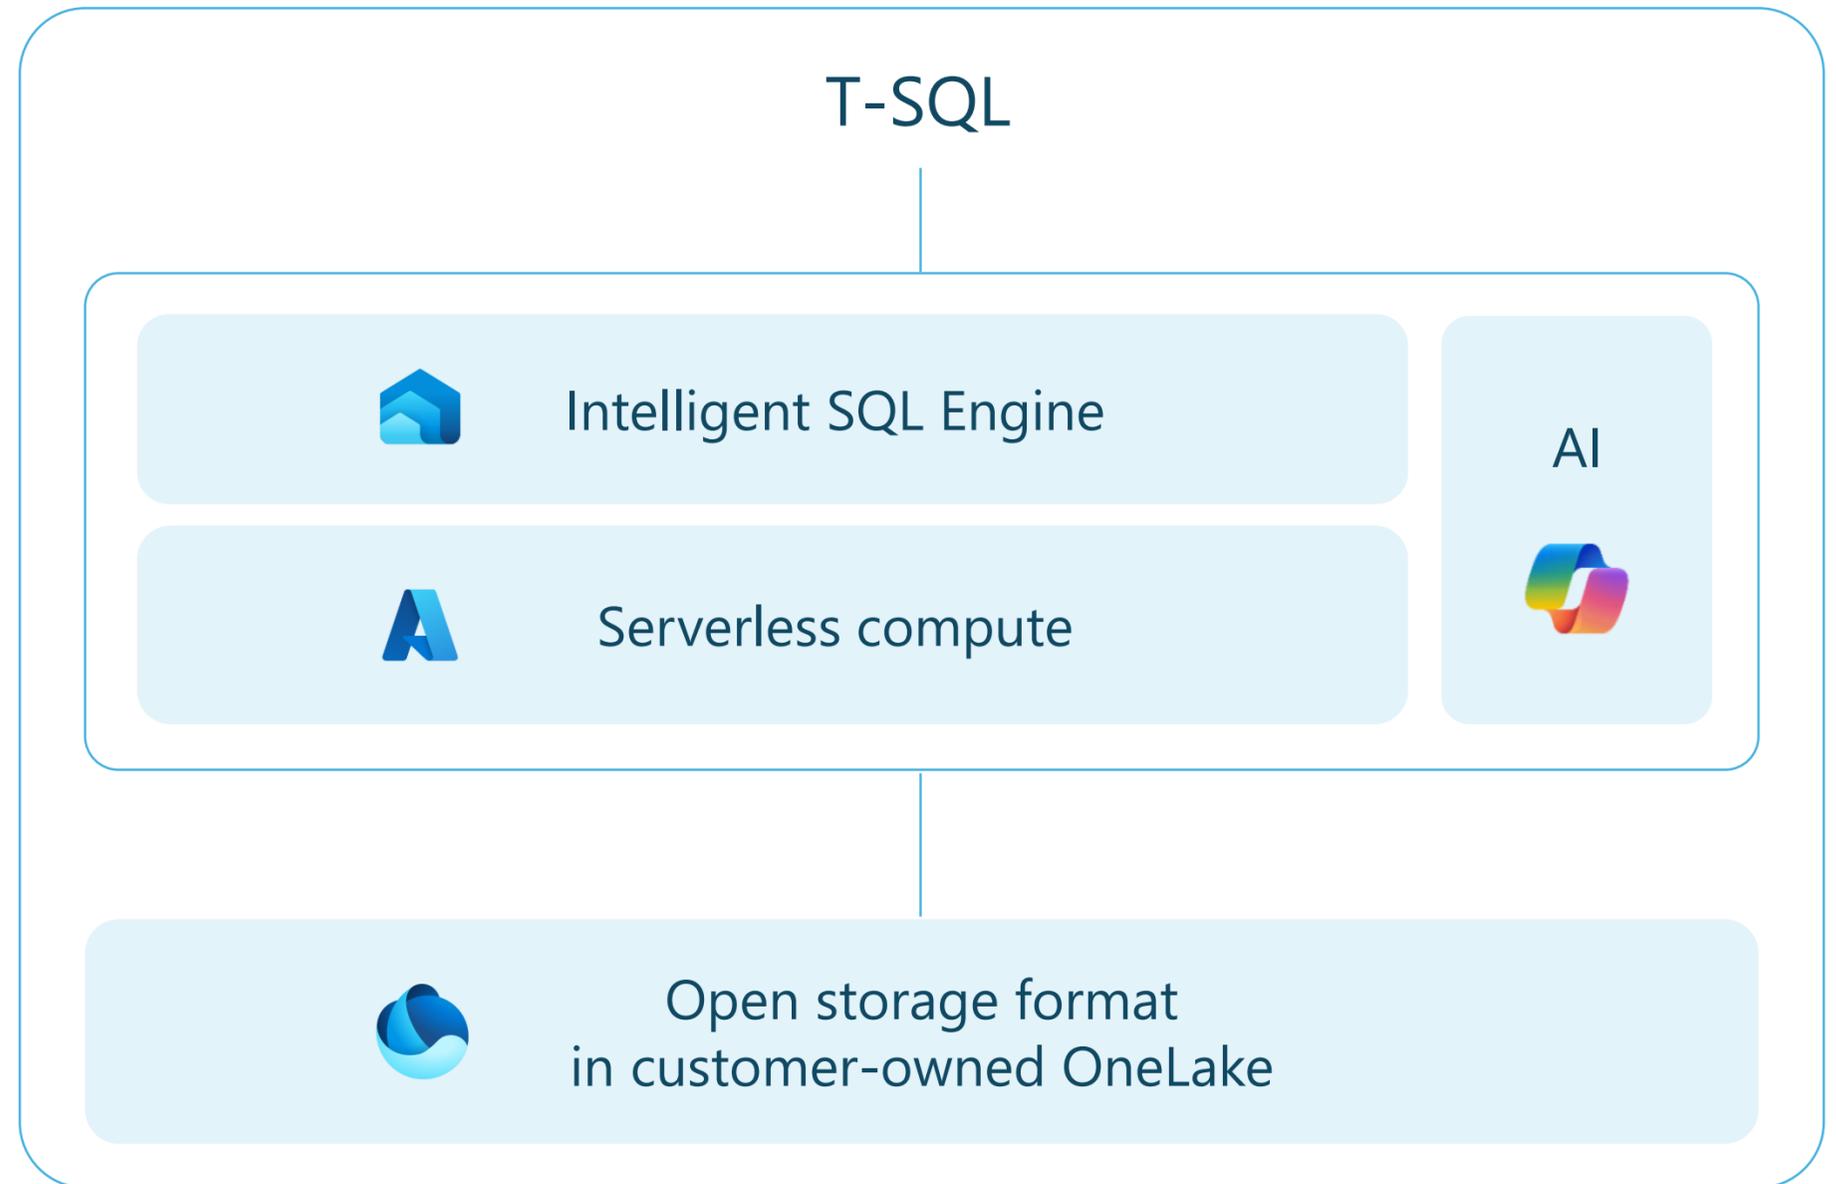
Interactive BI reporting with Power BI and Direct Lake

Agentic AI as a first-class capability, enabling business users to **chat with data**

Secure and govern the most complex data on a unified SaaS platform with **fewer moving parts**.

Capacity-based cost model and quicker **scale-up/down**

Rich **T-SQL** surface area





Why migrate: Data Warehouse Performance

1 TB TPCH - 22 Queries - Warm Cache



Source: Based on benchmarking study conducted by Microsoft Corporation August 2025 using TPC- H like metrics on 1 TB dataset. Runtime and Cost per query were calculated for Azure Synapse Dedicated Pools DWU1000 and Fabric Data Warehouse F64.

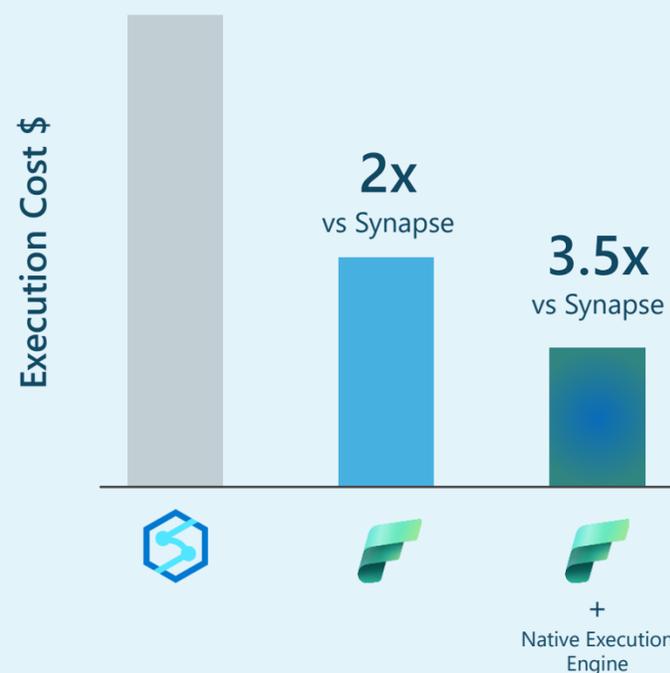


Why migrate: Fabric Data Engineering vs Synapse

Cost efficiency



TPC-DC 1TB

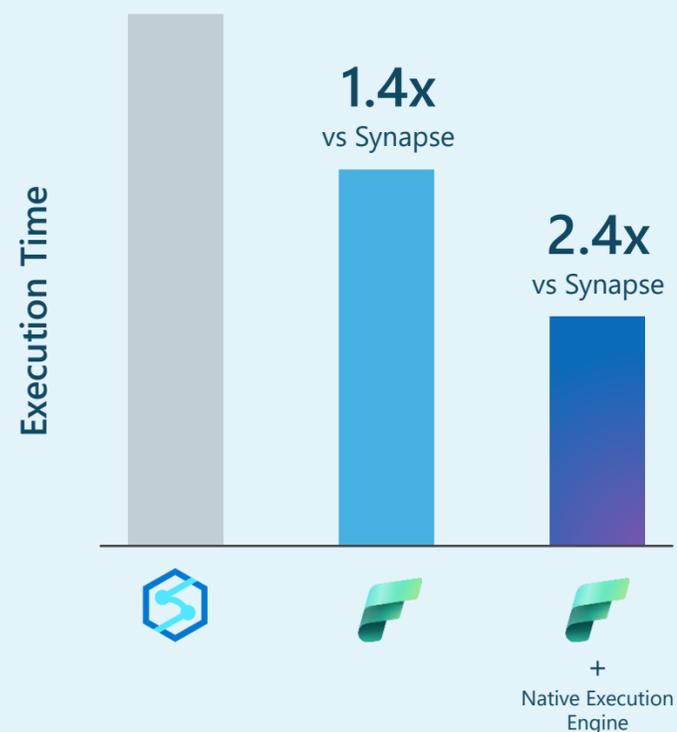


*Synapse Spark in non-US regions typically uses the same VM SKU as Fabric so there would be less perf diff compared to Fabric w/o NEE

Superior performance



TPC-DC 1TB



*Synapse Spark in non-US regions typically uses the same VM SKU as Fabric so there would be less perf diff compared to Fabric w/o NEE

SaaS-based Spark offering

Built-in starter pool requiring minimal Spark expertise

Enhanced Spark runtime with Microsoft innovations (Native engine, high concurrency, etc.)

Cross-workspace Lakehouse on OneLake using open Delta Lake format

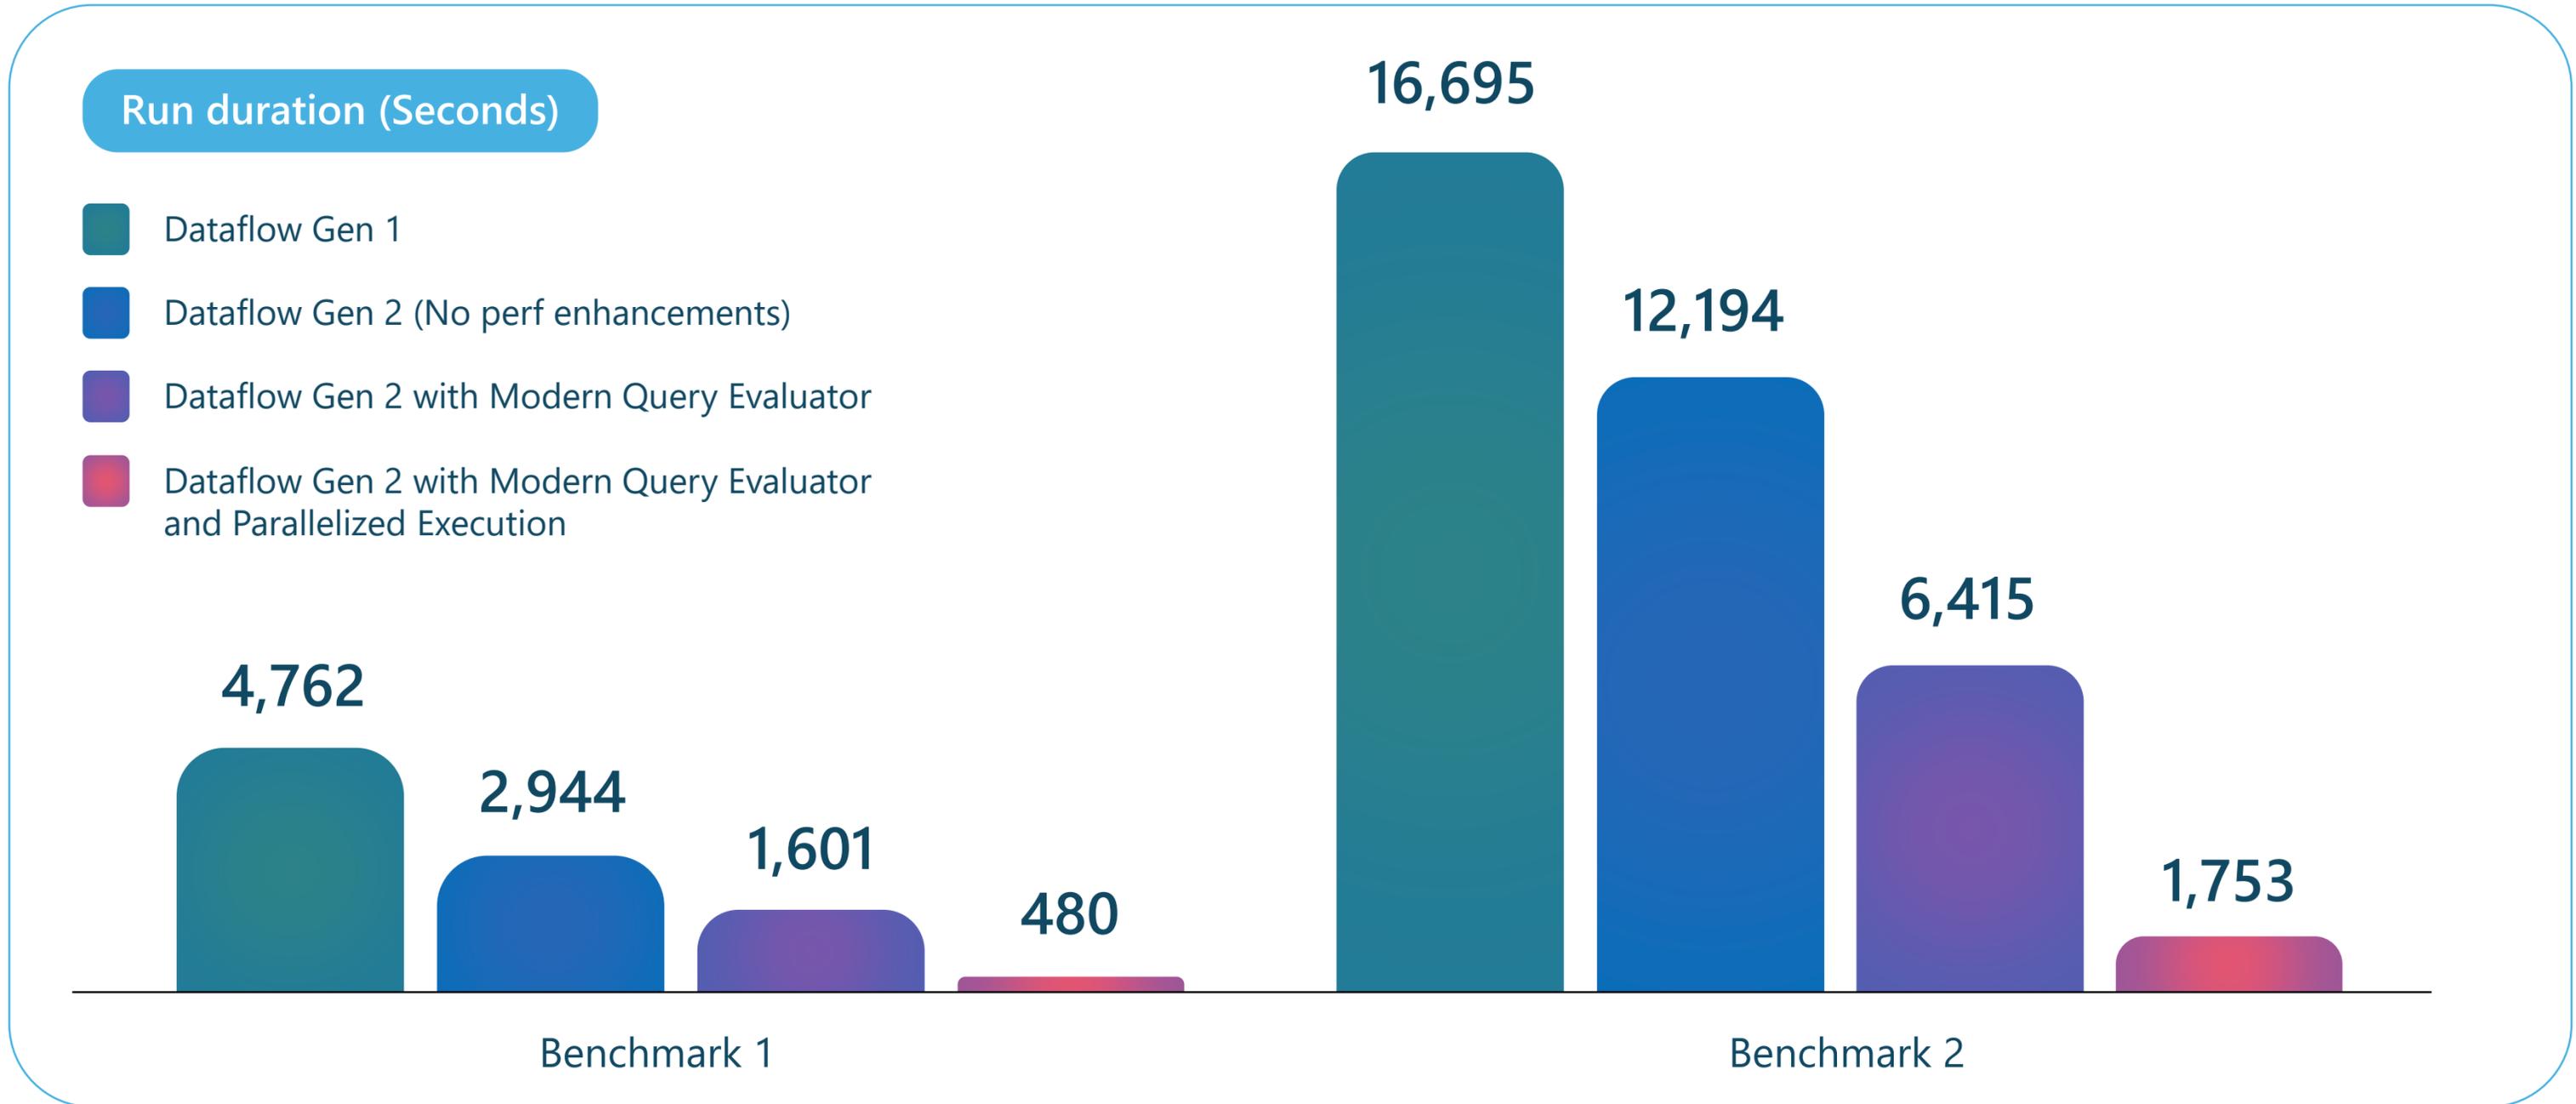
Integrated with Fabric data agents for natural language query and the

OneLake Catalog for governance and discoverability

Copilot-assisted capabilities for autonomous optimization and diagnostics



Why Migrate: Fabric Data Factory performance boost



Benchmark #1: NYCTaxi-2017, 1 partition, 9.8GB, 113M rows, 2 added columns
Benchmark #2: NYCTaxi-20xx, 5 partitions, 32.2GB, 365M rows, 2 added columns



Migration Tools



Data Warehouse

DW Migration Assistant (GA)

- Live connectivity (coming soon)



Data Engineering

Discovery and Assessment Tool

Migration Assistant (Public Preview)

NEW

Semi-auto Migration Script



Data Factory

PowerShell Upgrade Script

Migration Assessment Assistant

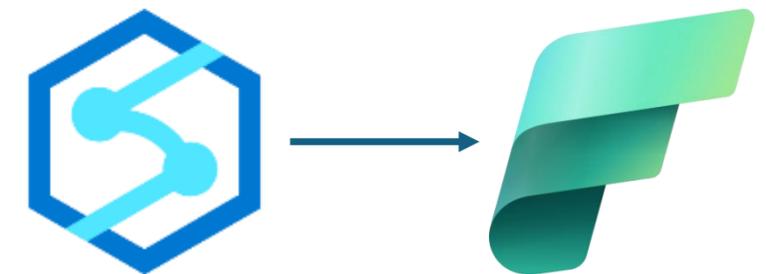
Fabric Data Factory Migration Assistant
(Microsoft CAT fabric-toolbox)



Cloud Accelerate Factory

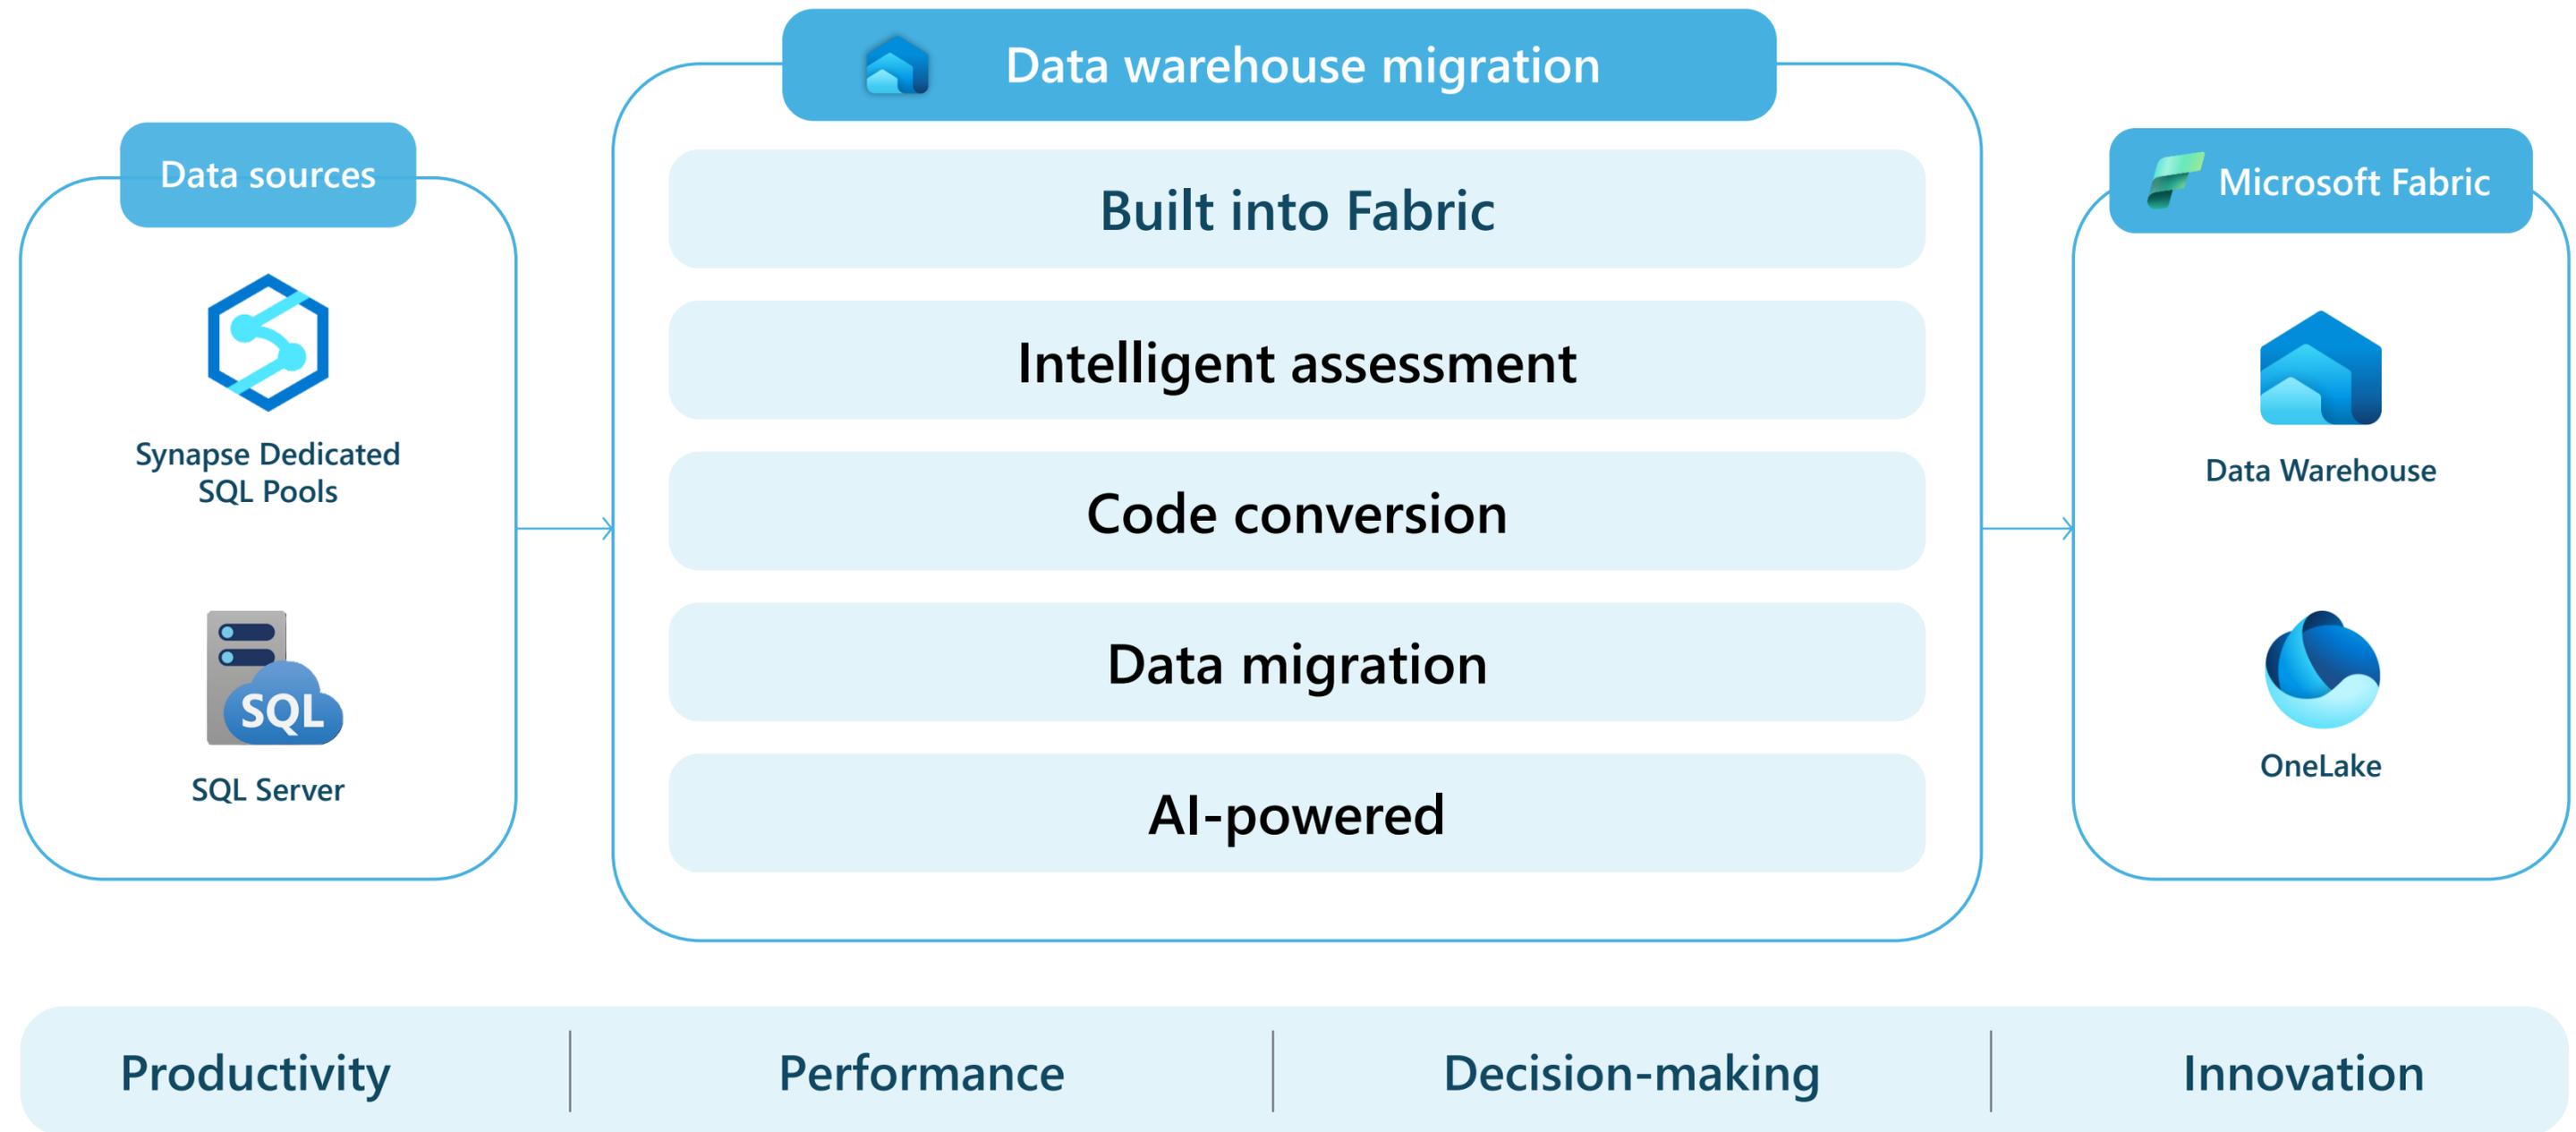


3rd Party tools



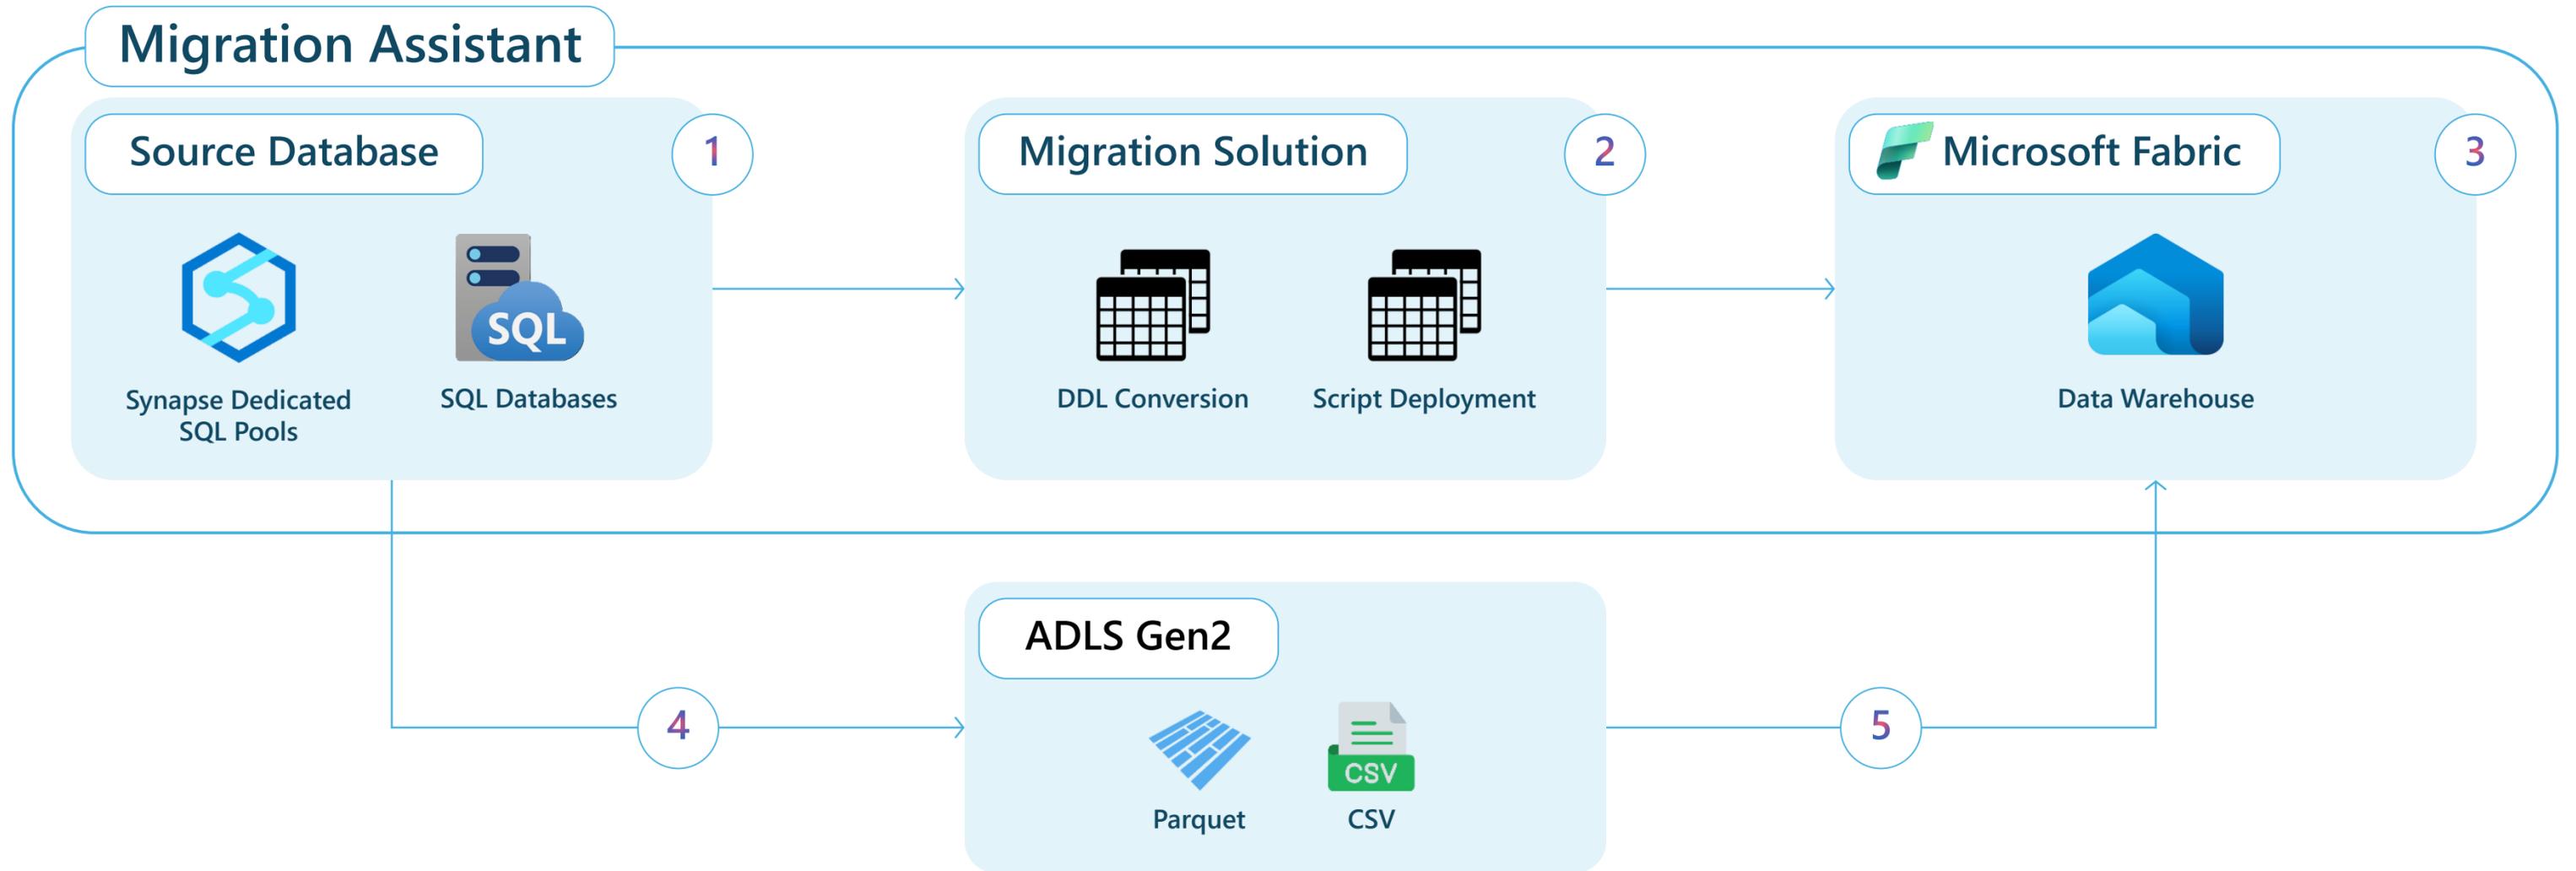


Data Warehouse Migration Assistant

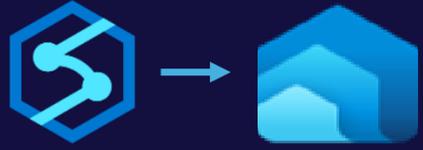




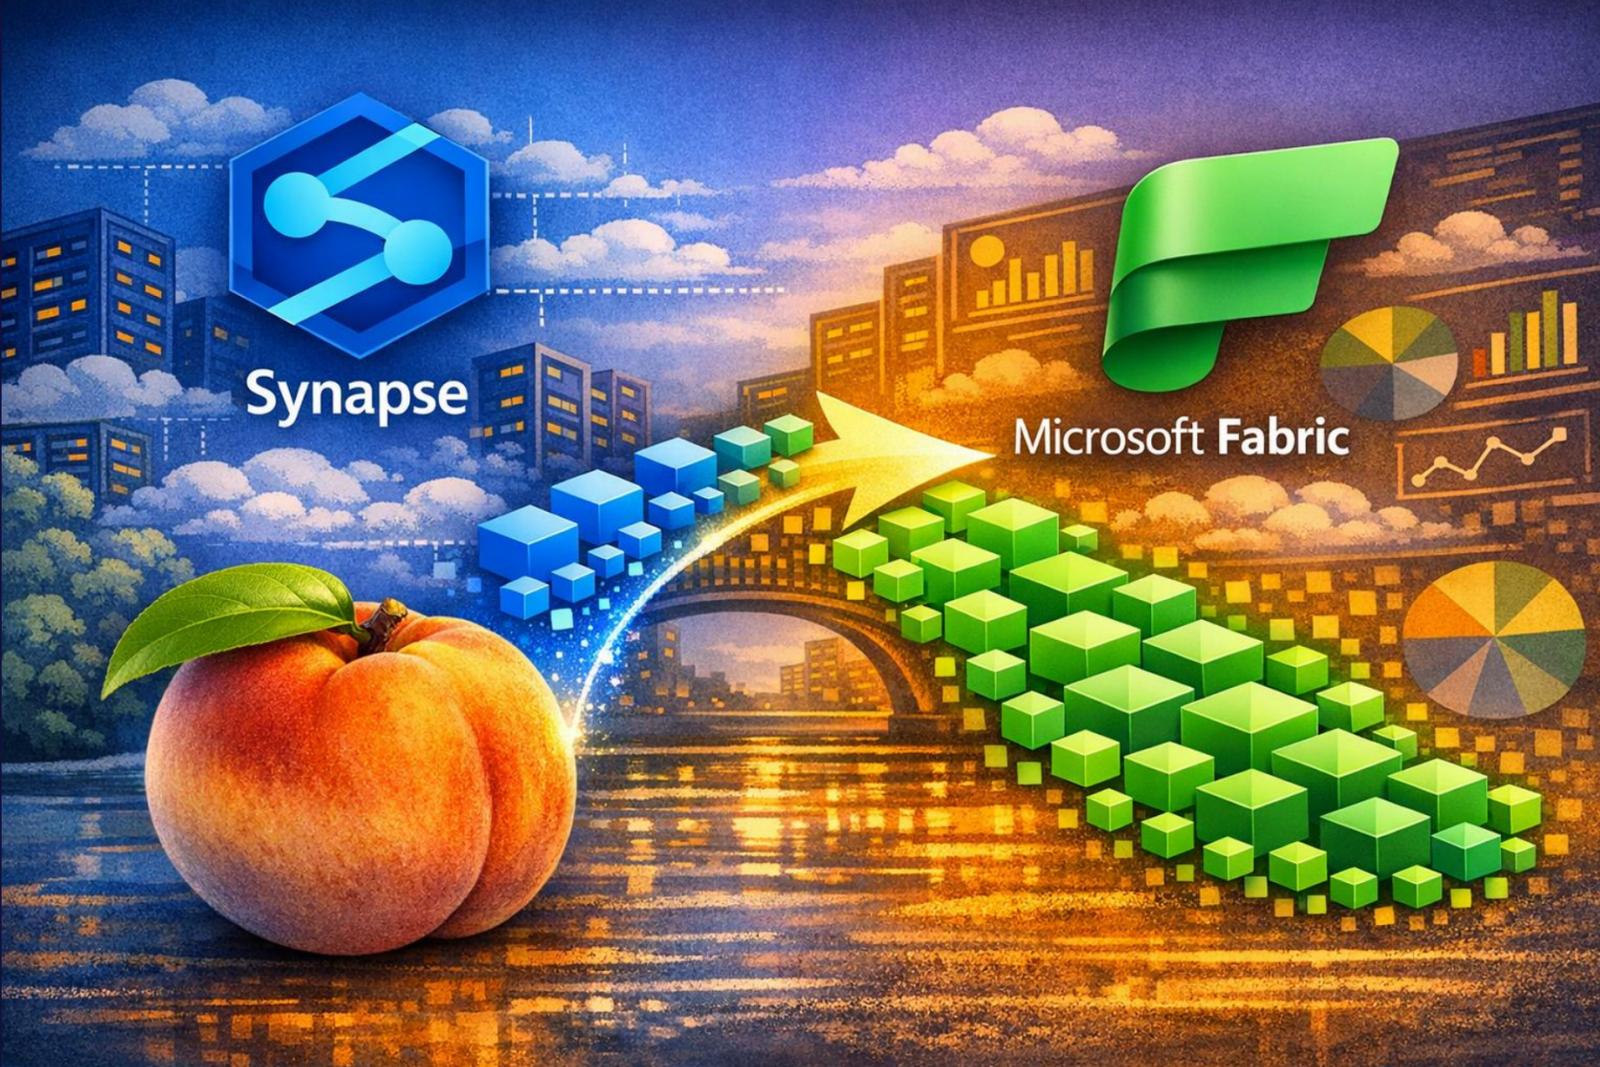
Migrating to Fabric Data Warehouse



- 1 Identify a source and create a DACPAC file
- 2 Migrate the data warehouse using the DACPAC file
- 3 Review the successfully migrated objects in the Fabric data warehouse
- 4 Copy the source data to ADLS Gen2 into Parquet or CSV format
- 5 Load the data to the Fabric data warehouse using the COPY INTO T-SQL command or your choice of tool (ADF, Pipeline, Copy Job etc)



DEMO
Data Warehouse
Migration Assistant





Choose from predesigned task flows or add a task to build one
Select from one of Microsoft's predesigned task flows or add a task to start building one yourself.

Select a predesigned task flow Add a task ▾

→ Import a task flow



There's nothing here yet
Add something new, or upload something to see them here.



Warehouse Migration Assistant Considerations

AI features require Copilot to be activated

- Requires a paid capacity (F2 or higher) - Copilot is not supported in trial capacities

Copy metadata and data

- DACPAC is (currently) required to copy metadata
- Copy Jobs are used to copy the data. For large data copies:
 - Check your VNET Data Gateway (if using one) for bottlenecks
 - Use CETAS to export the data to ADLS Gen2 then COPY INTO to import data into Fabric
 - [Efficiently Moving Data from Dedicated SQL Pool to Fabric Data Warehouse | by Mark Pryce-Maher | Medium](#)

Warehouse collation

- Migration Assistant creates the new warehouse with a case-insensitive collation
- Once a warehouse or SQL analytics endpoint is created, the collation can't be changed

Review “State” of migrated objects

- “Adjusted” means metadata was changed during the translation
 - Datatypes or T-SQL constructs are changed to ones supported in Fabric Data Warehouse

Review failed objects

- Copilot can suggest fixes, but manual input will be required



Other Migration Considerations

aka.ms/fabricroadmap

T-SQL incompatibilities

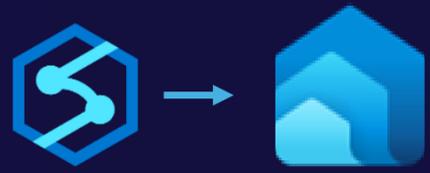
- IDENTITY column differences
- Datetime isn't supported
- HASHBYTES should use SHA2 instead of SHA1

ISOLATION LEVEL

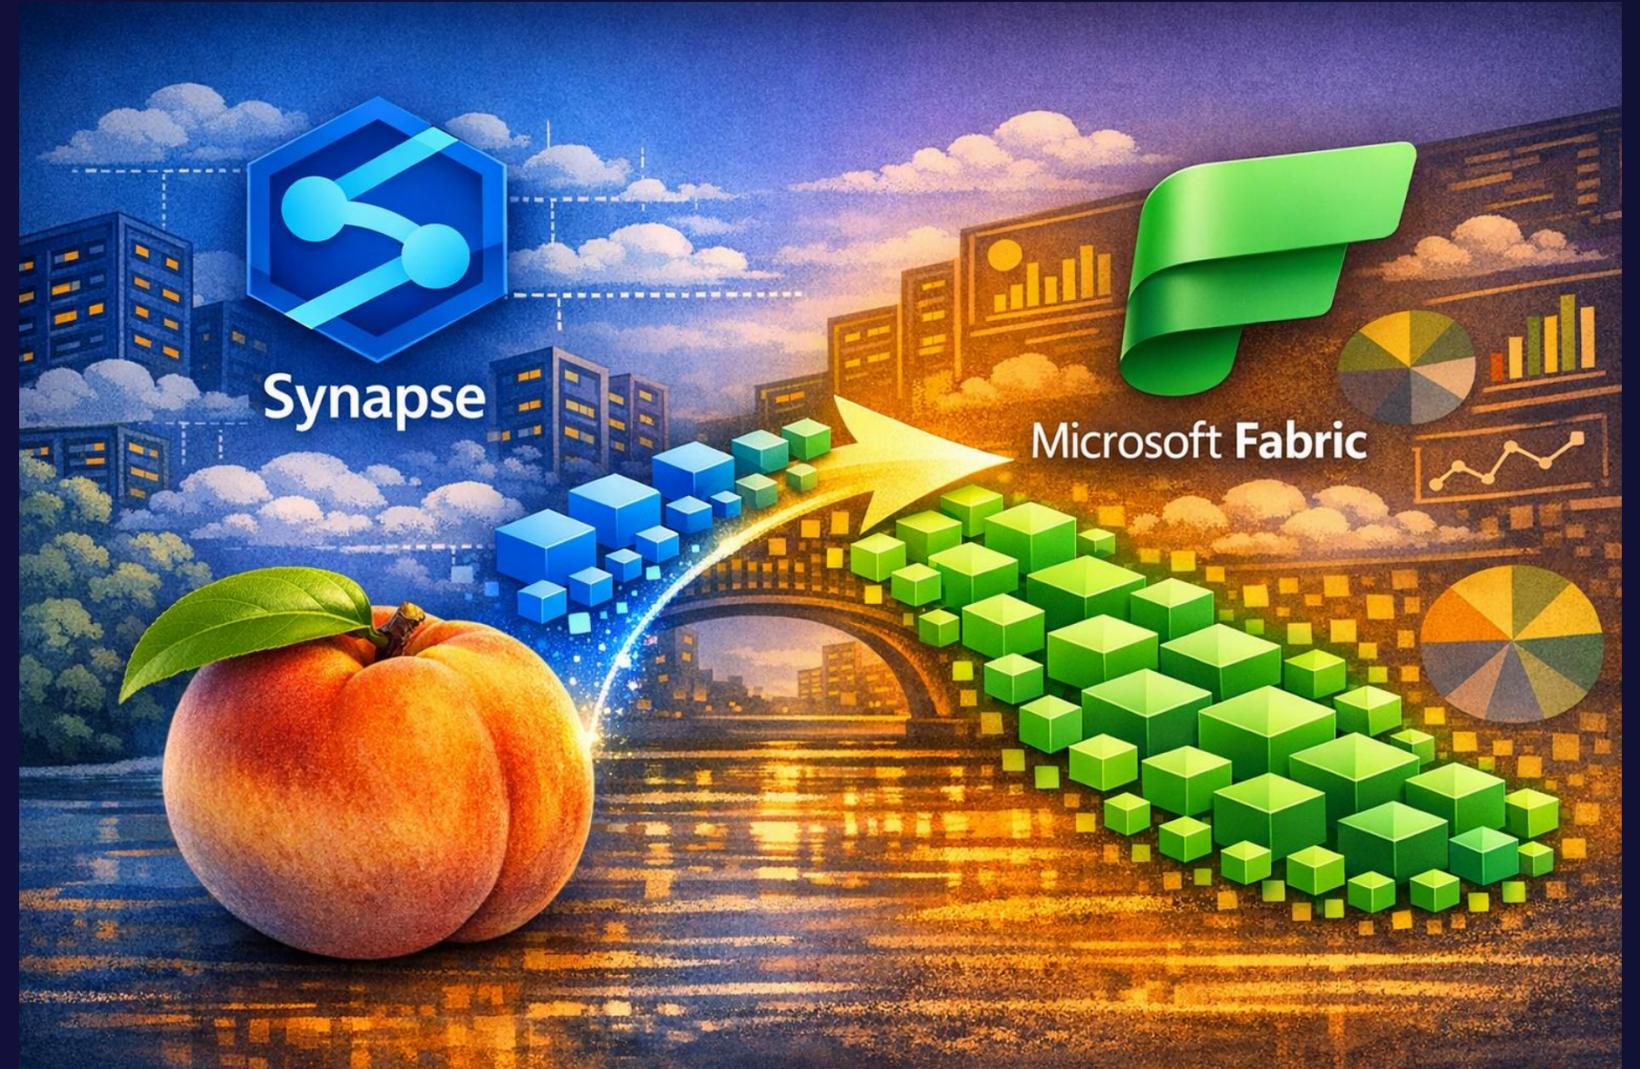
- Fabric Warehouse supports SNAPSHOT isolation only

Run an assessment of the source code

- Assess T-SQL code before migration to identify issues and anti-patterns
- Fix as many of these prior to migration as possible



DEMO Fabric Warehouse Code Compatibility Assessment





Run T-SQL Code Compatibility Assessment

Upload a SQL project zip (or .sqlproj) to generate Fabric compatibility reports.

Choose File

No file chosen

Upload and Assess

Recent Jobs

Assessment	Created	Reports	Open	Download All
Assessment result 2026-03-11 21:32	2026-03-11 21:32:32	21	Open summary	Download zip



What about Synapse Serverless SQL pools?

Migration Assistant only supports Dedicated SQL pools

- For Serverless Pools, assess external tables and views by source format (Delta, Parquet, CSV)

Option 1 - Keep **ADLS Gen2** as default storage and use shortcuts

- Create **OneLake shortcuts** to your existing ADLS Gen2 paths
 - For Delta, create the shortcut under **/Tables**
 - Shortcuts for **CSV/JSON/Parquet created** in **/Files**
 - **N.B.** These **/Files** shortcuts won't be automatically discovered by the SQL endpoint

Option 2 - Move data to **OneLake** as default storage

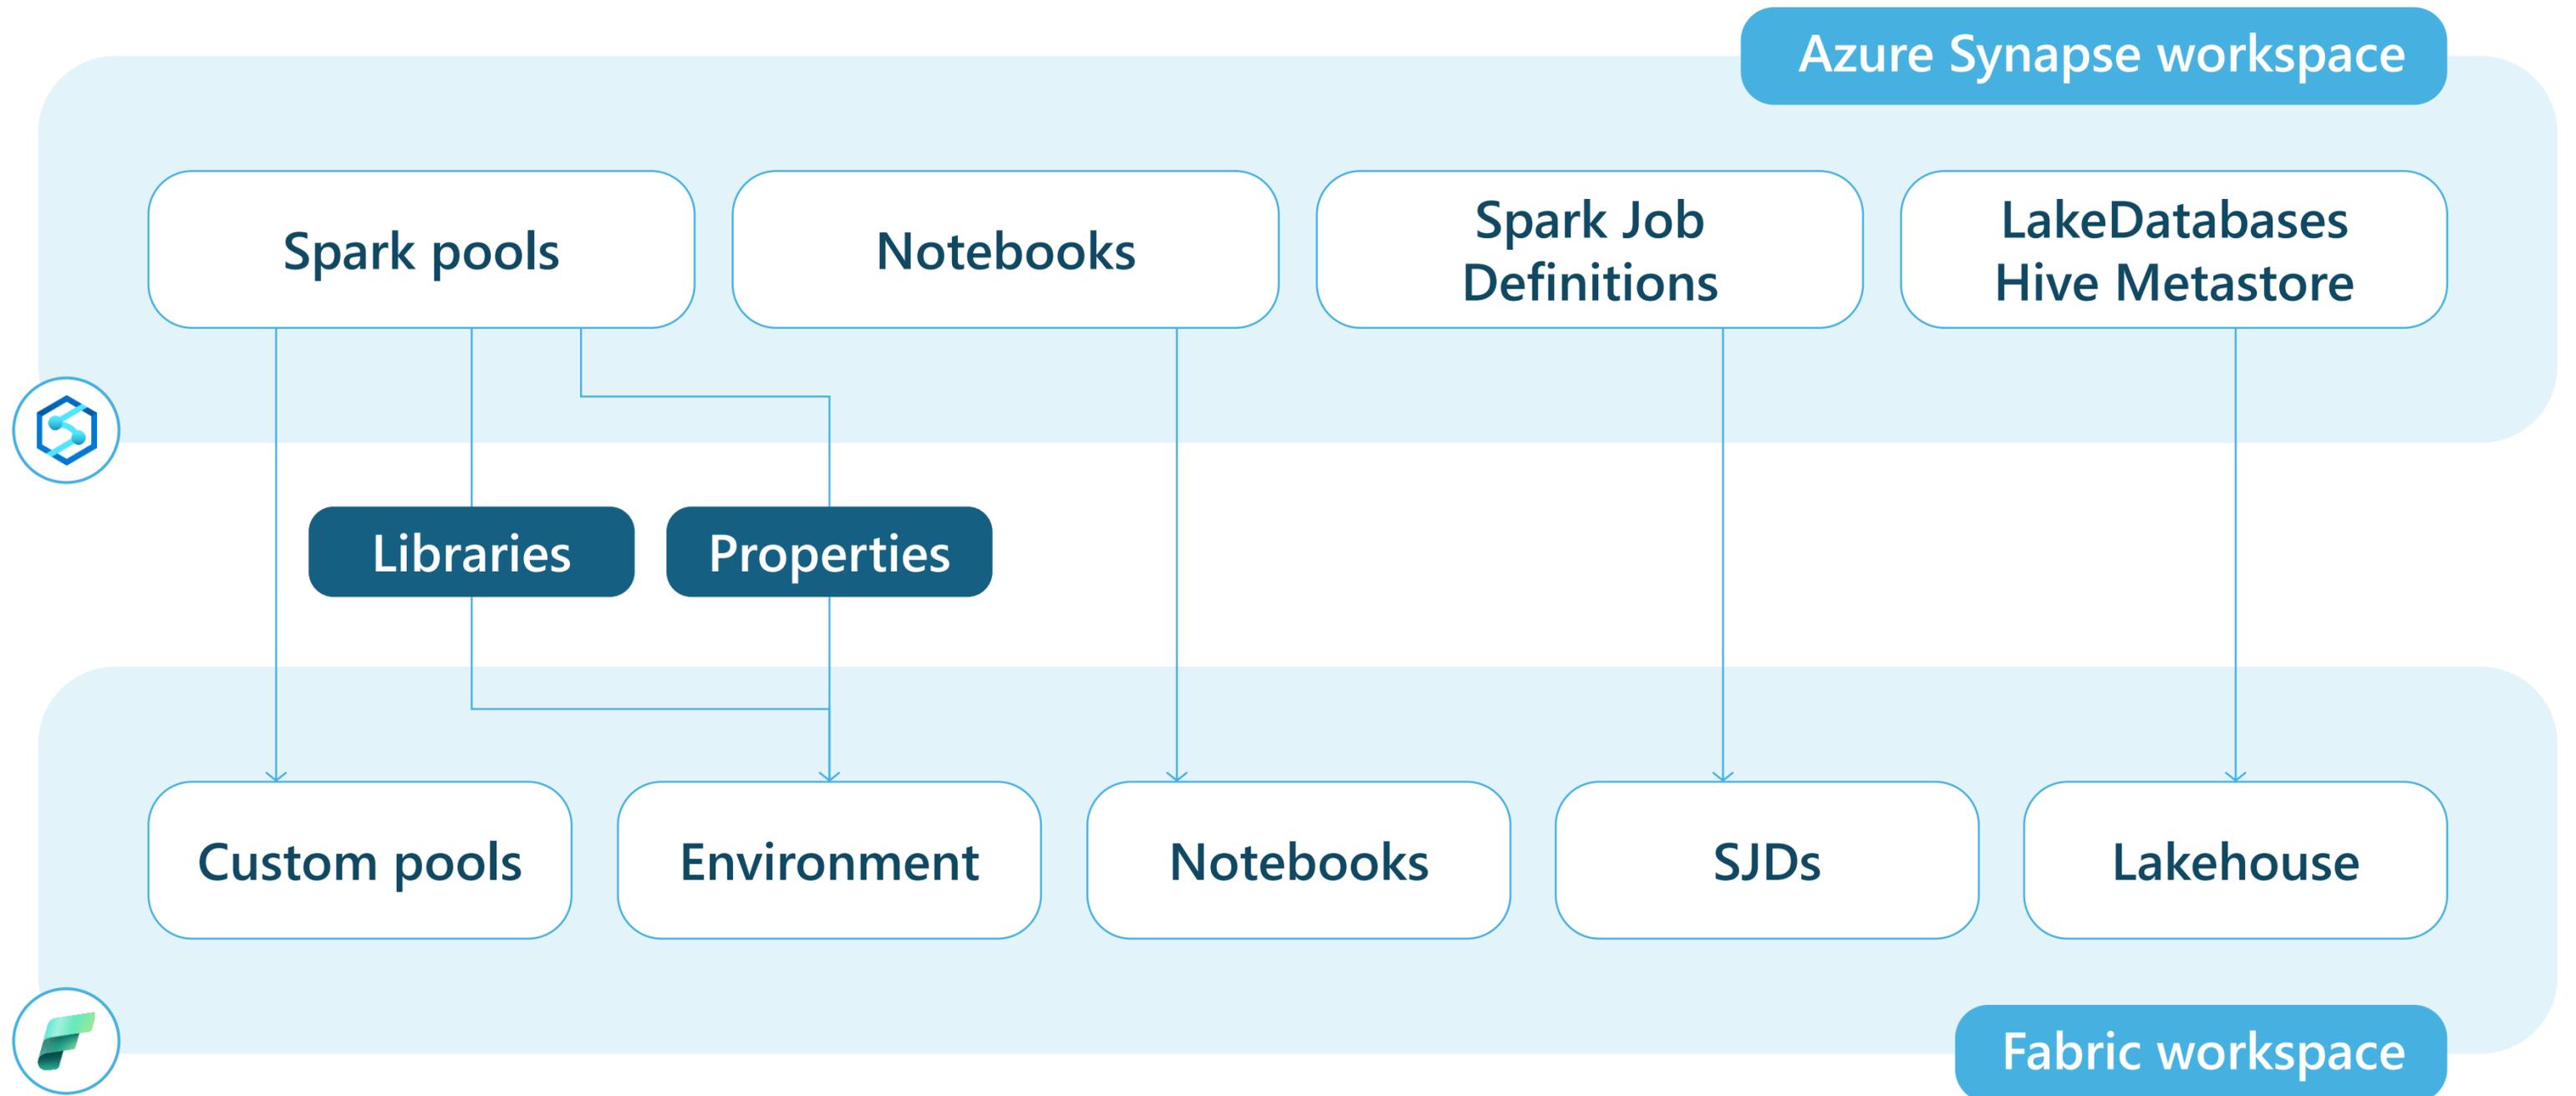
- Migrating into OneLake reduces external dependencies, simplifies governance

Do you even need a warehouse?

- Start with *Lakehouse + SQL endpoint* (plus views) to replace Serverless querying.
- Add a **Warehouse** only when you need writable T-SQL or Warehouse-specific features as the SQL Analytics endpoint is **read-only**



Data Engineering migration





Data Engineering items migration

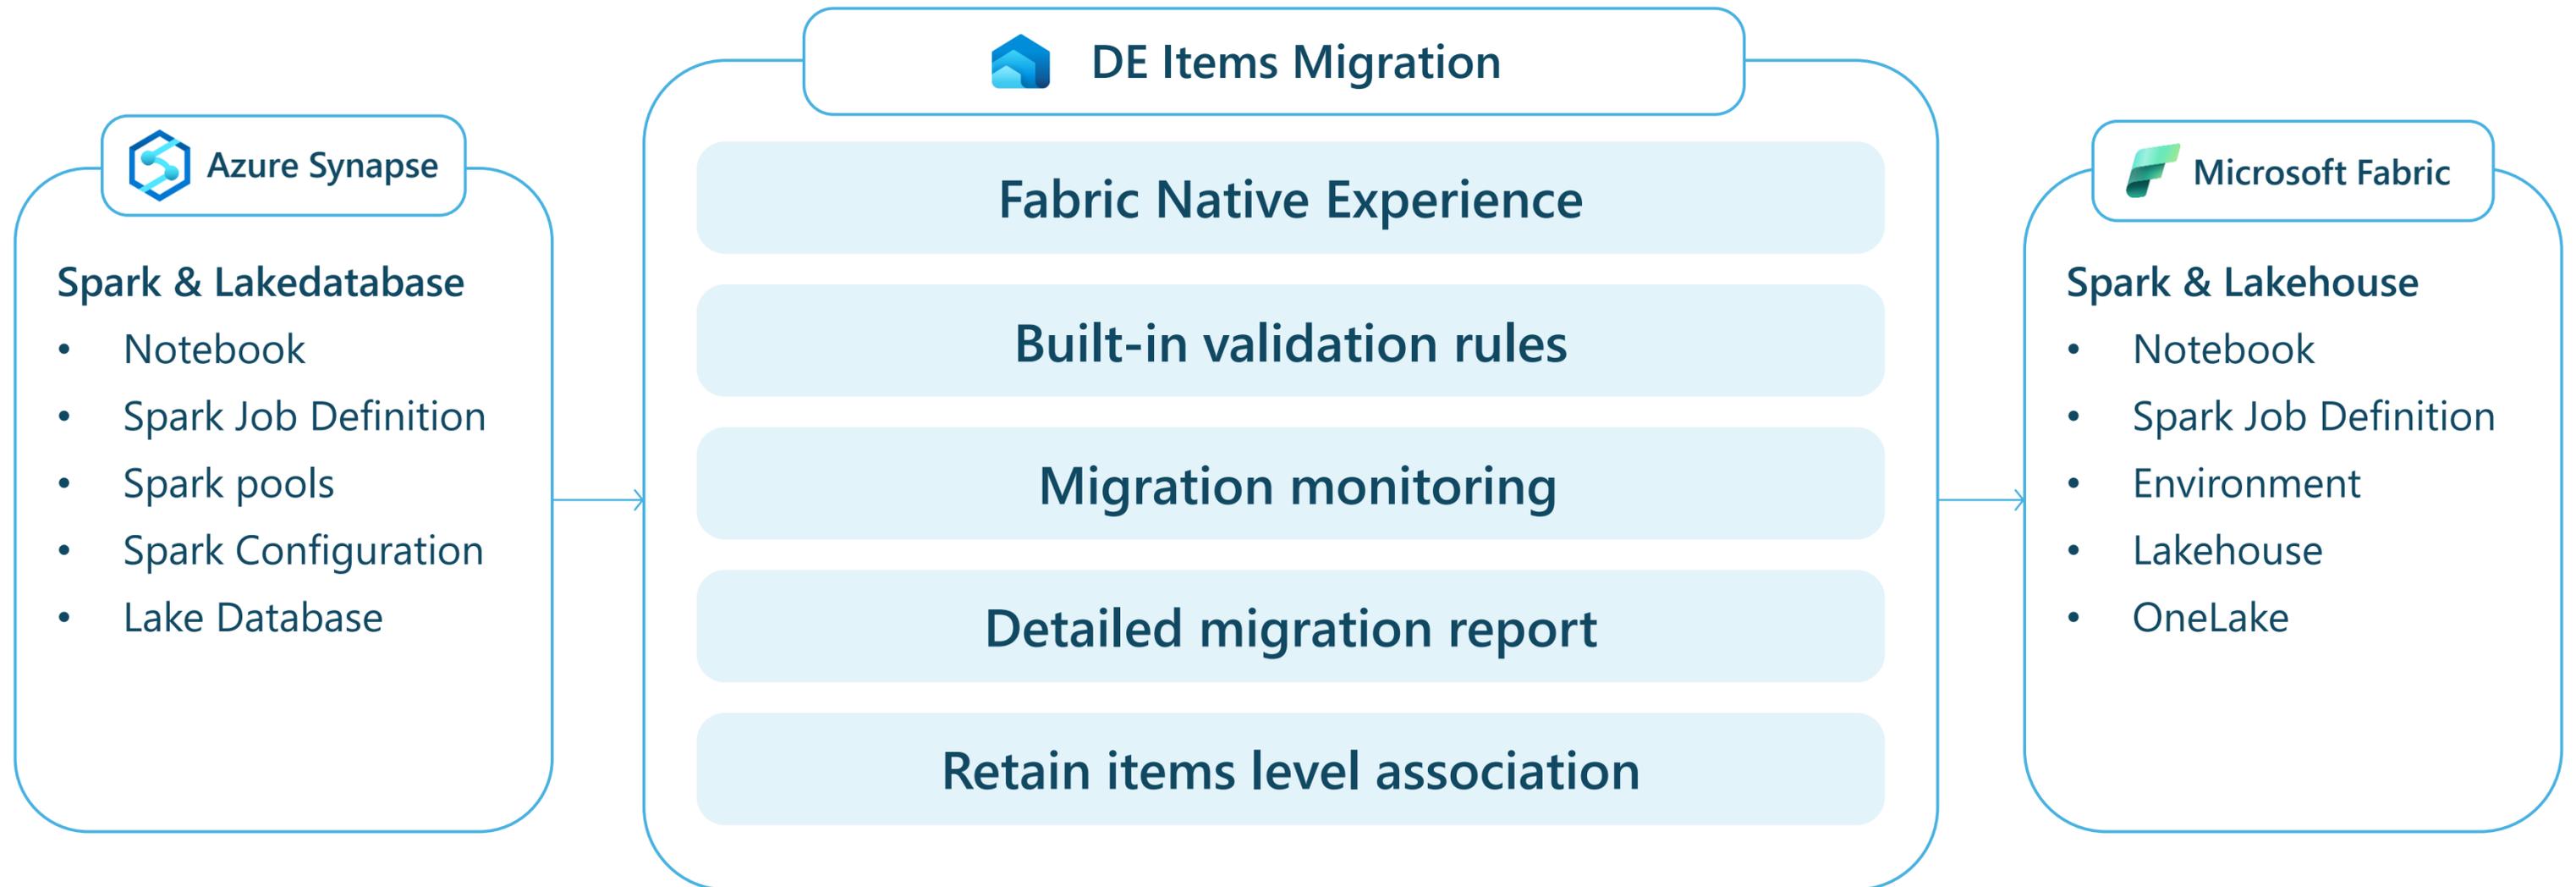
 Azure Synapse

 Fabric Spark

	Spark pools	➔ Starter and Custom Spark pools
Compute	Spark configurations (pool level)	➔ Environments
	Libraries (pool level)	➔ Environments
Resources	Notebooks (Python, Scala, Spark SQL, R, .NET)	➔ Notebooks (Python, Scala, Spark SQL, R)
	Spark job definitions (Python, Scala, .NET)	➔ Spark job definitions (Python, Scala, R)
	Notebook pipeline activity	Fabric pipeline Notebook activity
	Spark Job Definition pipeline activity	Fabric pipeline SJD activity
Data	ADLS Gen2 as main storage	<ul style="list-style-type: none"> • OneLake as main storage • Use shortcuts to ADLS Gen2
HMS	Internal	➔ <ul style="list-style-type: none"> • Use shortcuts > Delta tables only • Lakehouse /Tables > Delta & non-Delta tables
	External (Azure SQL DB based)	

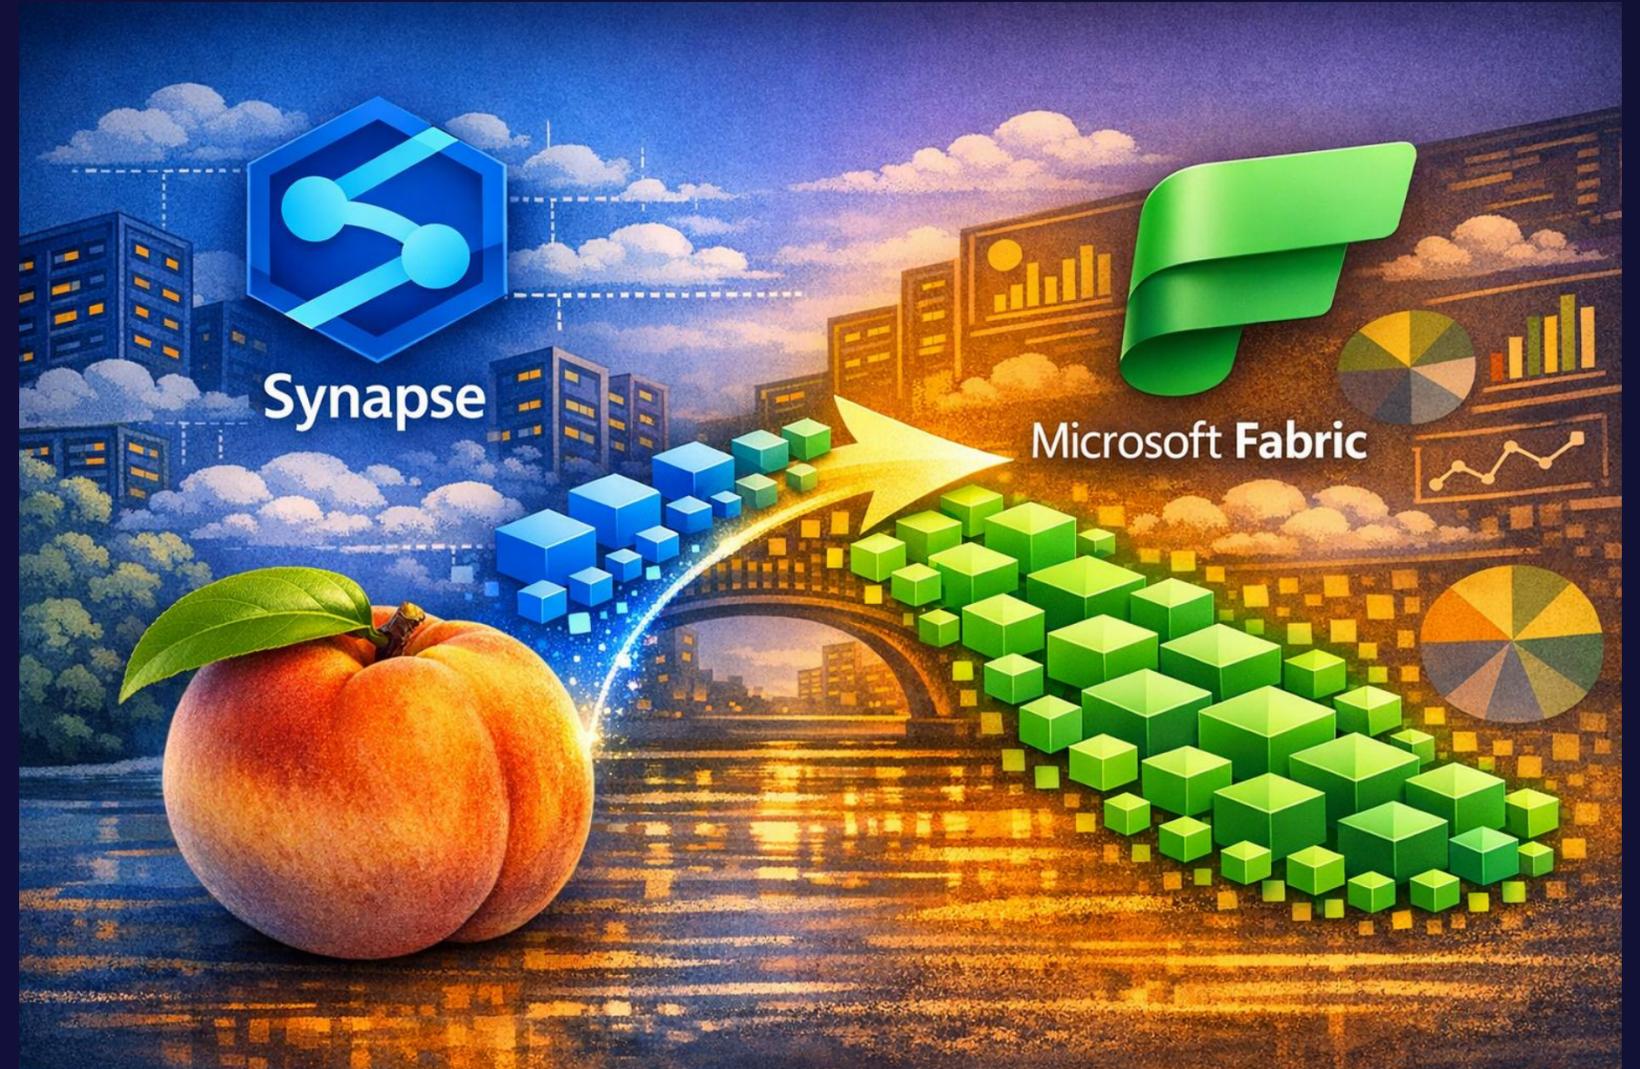


Data Engineering Migration Assistant (Public Preview) **NEW**





DEMO Data Engineering Migration Assistant



Welcome to Fabric

Create a workspace with a predesigned template called a task flow. Task flows keep your items organized. [Learn more](#)

<
>



New workspace



General



Basic data analytics



Data analytics using a SQL ...



Medallion



Event analytics



Lambda



Sensitive data insights



Basic machine learning models



Even

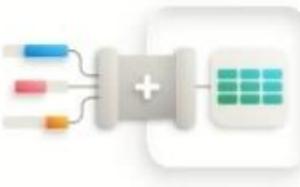
Learn more about Fabric

<
>

What is Microsoft Fabric?
Watch this 1-minute introductory video



Ingest data into Fabric
Complete an end-to-end tutorial for Data Factory



Build a lakehouse
Complete an end-to-end tutorial for Lakehouse



Build a warehouse
Complete an end-to-end tutorial in Data Warehouse



^ Show less

Quick access


Recent workspaces


Recent items


Favorites




Filter

	Name	Opened
	Synapse Migration Demo 	now
	My workspace  	now
	migration2  	12 days ago



How to Migrate: Additional Considerations

Spark Code & Runtime

- Replace mssparkutils with notebookutils and validate runtime versions
- Move Linked services and data source connections
 - Cloud connections, On-Prem Data Gateway, VNet Data Gateway
- Synapse ADLS Gen2 URLs (abfss) should be updated to read/write through OneLake or shortcuts
 - `spark.table("schema.table"), df.write.format("delta").saveAsTable("schema.table")`

Libraries & Dependencies

- Rebuild any custom libraries as Fabric Environments. Test any upgraded libraries for syntax changes

Security & Governance

- Assigning workspace RBAC roles and access policies
- (Managed) private endpoints and VNets
- Re-point Power BI, apps and reports to use Fabric endpoints

Validation

- Run workloads in parallel and validate schema/data.
- Run a code review and look for deprecated APIs

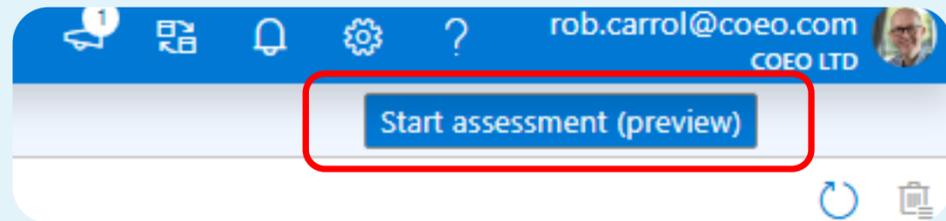


Data Factory Migration Considerations

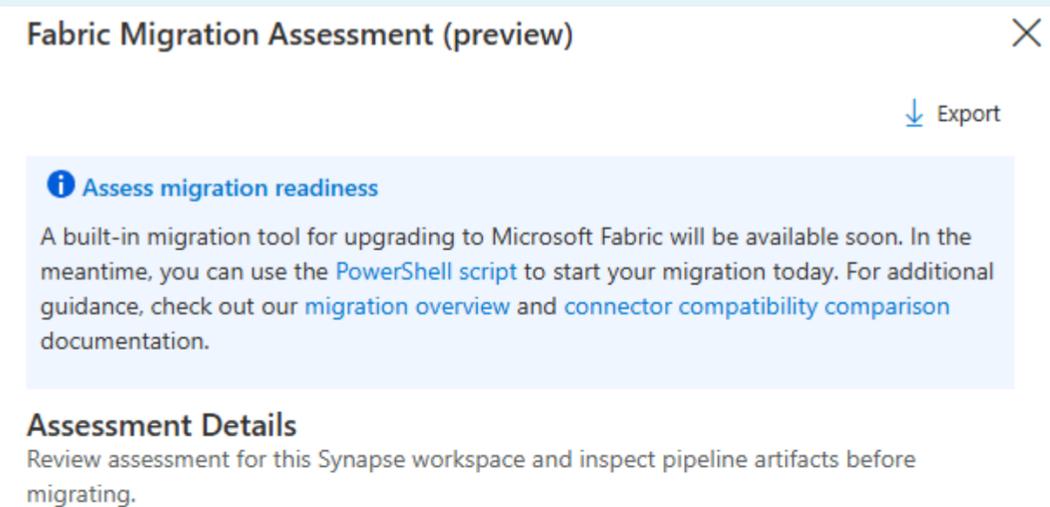
Feature	Fabric vs ADF
Integration runtimes	Fabric uses cloud-based compute by default; ADF requires configured IRs.
Pipelines	Fabric includes SaaS-based activities and uses different JSON definitions.
Linked services	Fabric replaces Linked Services with in-activity Connections.
Datasets	Fabric removes datasets - data properties are defined inline.
Dataflows	Fabric uses Power Query; ADF Data Flows use a different engine & language.
Triggers	Fabric uses Activator framework vs ADF standalone triggers.
Debugging	Fabric removes separate debug mode for simplified authoring.
Change Data Capture	Fabric uses Copy jobs instead of CDC artifacts.
Azure Synapse Link	Mirroring replaces Synapse Link for replication.
SSIS	Fabric doesn't support SSIS IRs but can invoke ADF for SSIS execution.
Invoke pipeline	Fabric enhances Execute pipeline with cross-platform invocation.



Pipeline migration readiness assessment in Synapse



Click "Start assessment (preview)" to run a migration readiness assessment of your Synapse pipelines.



- Assess your Synapse pipelines for "Fabric readiness".
- Expand results to view details of assessment
- Export results to plan your migration.

Migration Summary

Status	Pipelines
● Ready	60 of 61
● Needs review	0 of 61
● Coming soon	1 of 61
● Not compatible	0 of 61

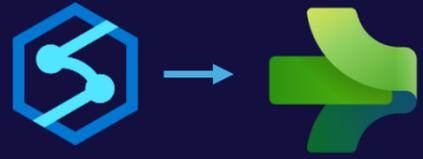
> ● Copy Files Recursively by LastModified sKey dKey v1

▼ ● Copy Parquet to Delta sKey dKey v1

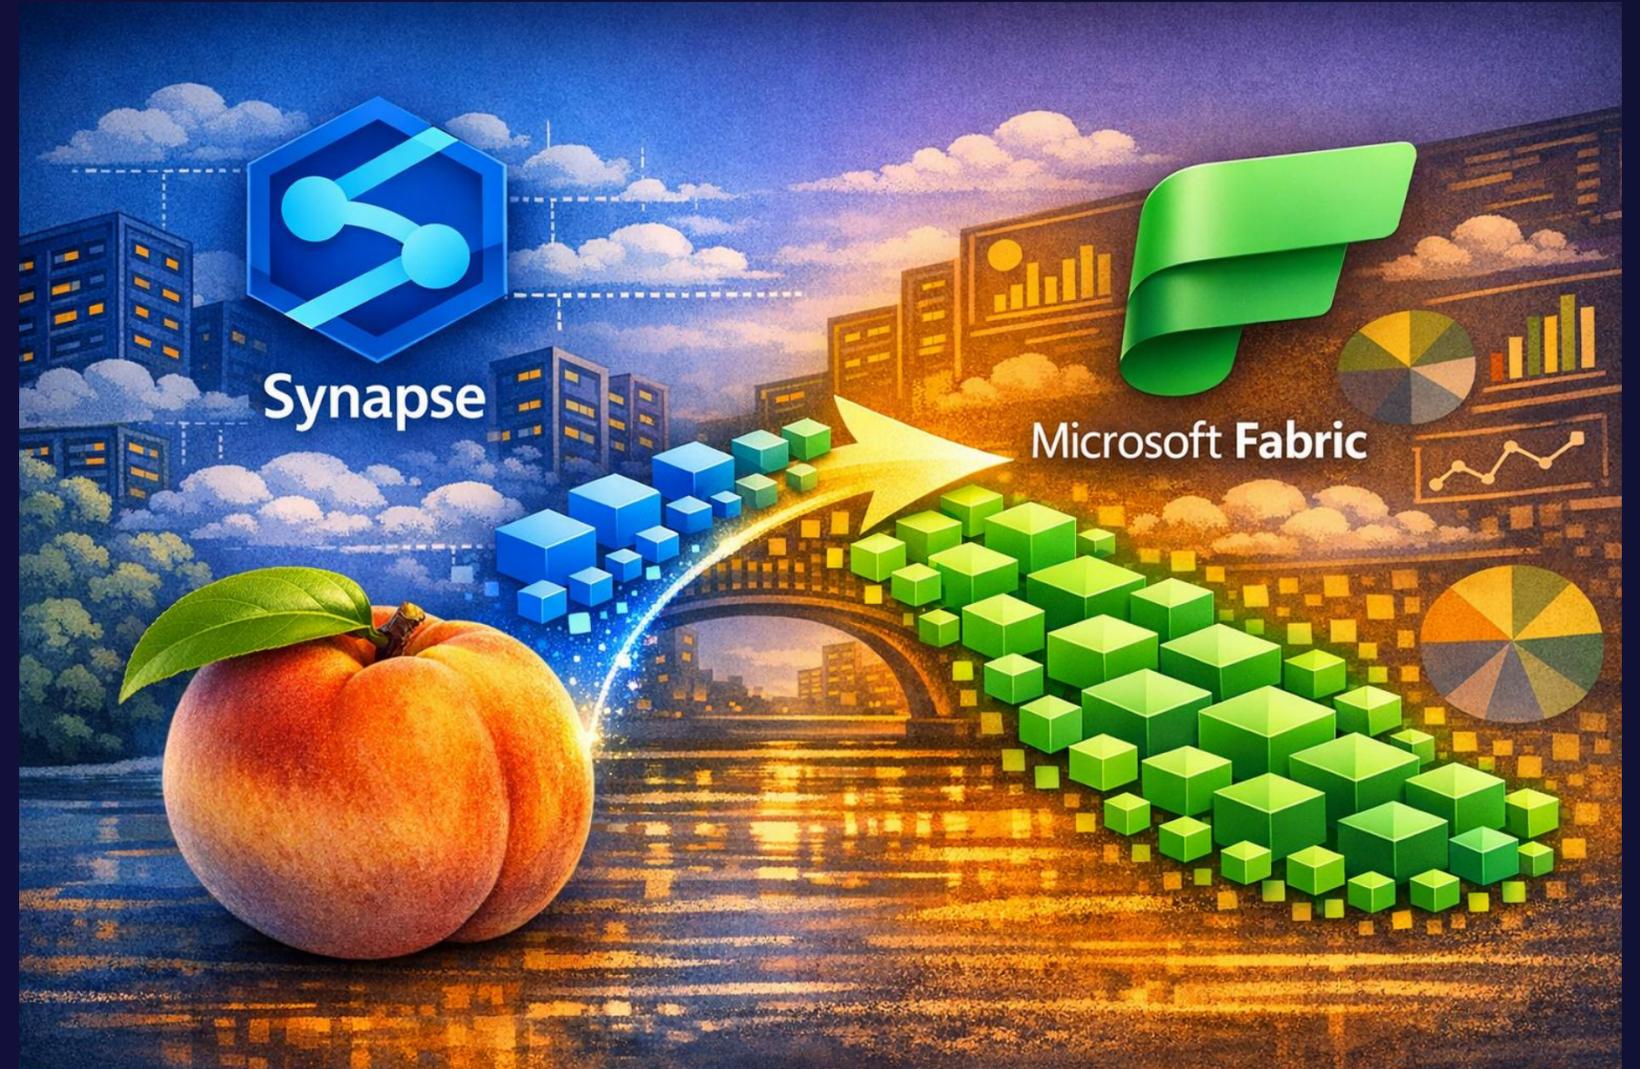
Name	Type	Status
Transform_Data	ExecuteDataFlow	● Coming soon
Log Pipeline Start	AzureFunctionActivity	● Ready
Log Pipeline End	AzureFunctionActivity	● Ready

> ● Copy SFTP Recursively by LastModified sBasic dMan v1

This resource is 'Coming Soon' because it depends on 'ExecuteDataFlow' with that same status.



DEMO Migration Assessment Tool





Synapse Analytics workspace synodince1ddev

New ▾



 **Ingest**
Perform a one-time or scheduled data load.

 **Explore and analyze** [↗](#)
Learn how to get insights from your data.

 **Visualize**
Build interactive reports with Power BI capabilities.

Discover more

 [Knowledge center](#)

 [Browse partners](#)

Recent resources

	Name	Last opened by you
	Initial Load - Static Tables	4 days ago
	Initial Load - Copy Static Tables from Raw to Curated	4 days ago

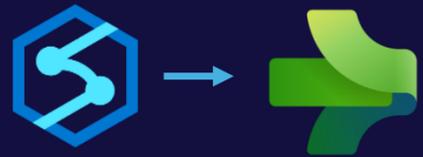
Feature showcase



SAP Knowledge Center

Azure Synapse has a comprehensive knowledge center with information on SAP data integration scenarios. This includes SAP connectors, templates, and more.

[Learn more](#) [↗](#)



DEMO Fabric Data Factory Migration Assistant



Files

main

Go to file

- .github
- .vscode
- accelerators
- media
- monitoring
- samples
- scripts
- tools
- .gitignore
- CODE_OF_CONDUCT.md
- LICENSE
- README.md
- SECURITY.md
- SUPPORT.md

fabric-toolbox / README.md

jdocampo fix readme link ✓ af6d116 · 2 months ago History

Preview Code Blame 83 lines (68 loc) · 5.31 KB Raw Download Edit



Welcome to the Fabric Toolbox

Fabric toolbox is a repository of tools, accelerators, scripts, and samples to help you work with Microsoft Fabric. This repository is brought to you by the Fabric Customer Advisory Team (CAT) and will continue to grow as we develop new tools and accelerators.

These assets should be treated as examples that you can use to create the solutions that are appropriate for your use case. If you have any issues, please use the [issues](#) tab of this repository and we will work to address issues on a best effort basis.

Monitoring

- **UPDATED!** [Fabric Cost Analysis \(FCA\)](#)
- **UPDATED!** [Fabric Unified Admin Monitoring \(FUAM\)](#)
- [Fabric Platform Monitoring \(Monitor Fabric with RTI and Capacity Events\)](#)
- [Workspace Monitoring Report Templates](#)
- [Fabric Spark Monitoring \(Monitor Spark with RTI\)](#)

Accelerators

- [BCDR Accelerator](#)
- [CICD: Git-Based Deployments](#)
- [CICD: Deploy using Fabric deployment pipelines](#)
- [CICD: Branch out to new workspace](#)

- 1 Upload ARM Template
- 2 Login & Authentication
- 3 Select Fabric Workspace
- 4 Managed Identity
- 5 Map Connections
- 6 Deploy Connections
- 7 Validation
- 8 Global Parameters
- 9 Component Mapping
- 10 Deployment
- 11 Complete

Upload ARM Template

Upload your ARM template to begin the migration process



Upload ARM Template

Drag and drop your ARM template JSON file here, or click to browse

[← Previous](#)

[Proceed with Migration →](#)



Pipeline migration tips

Run the assessment tool

- Plan for any incompatibilities, and pipelines that need to be manually migrated. For example, invoke Synapse notebook tasks, sFTP tasks

Make sure Synapse code is in a repo

- Download latest ARM template from your Synapse Repo (TemplateForWorkspace.json)

Create connections before migration

- Use the same names and the Fabric Data Factory migration assistant will map them automatically

Start small and test

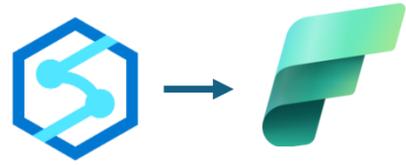
- Migrate a few pipelines at a time and test the output

Automate where possible

- Use the automated migration tools as much as possible
- AI agents were used to refactor some incompatible code

Use git integration and CI/CD

- Once migrated, enable git integration for your Fabric workspace and follow CI/CD release best practices for your pipelines in Fabric



Cloud Accelerate Factory

Jumpstart Azure projects with zero-cost deployment from Microsoft experts

Factory benefits

 **Get zero cost deployment assistance**

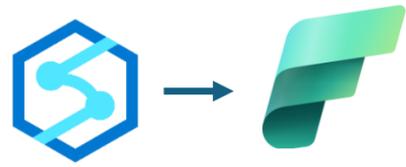
Get Microsoft assistance to deploy 30+ Azure services via joint delivery with an Azure partner no matter the project size

 **Accelerate with industry best practices**

Leverage proven strategies developed over thousands of customer engagements to complete deployments quickly

 **Unlock and scale your resources**

Prioritize your funding and skilled resources for the more advanced components or projects to boost business impact



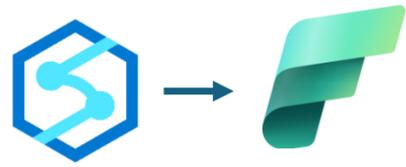
Third-party Fabric workloads



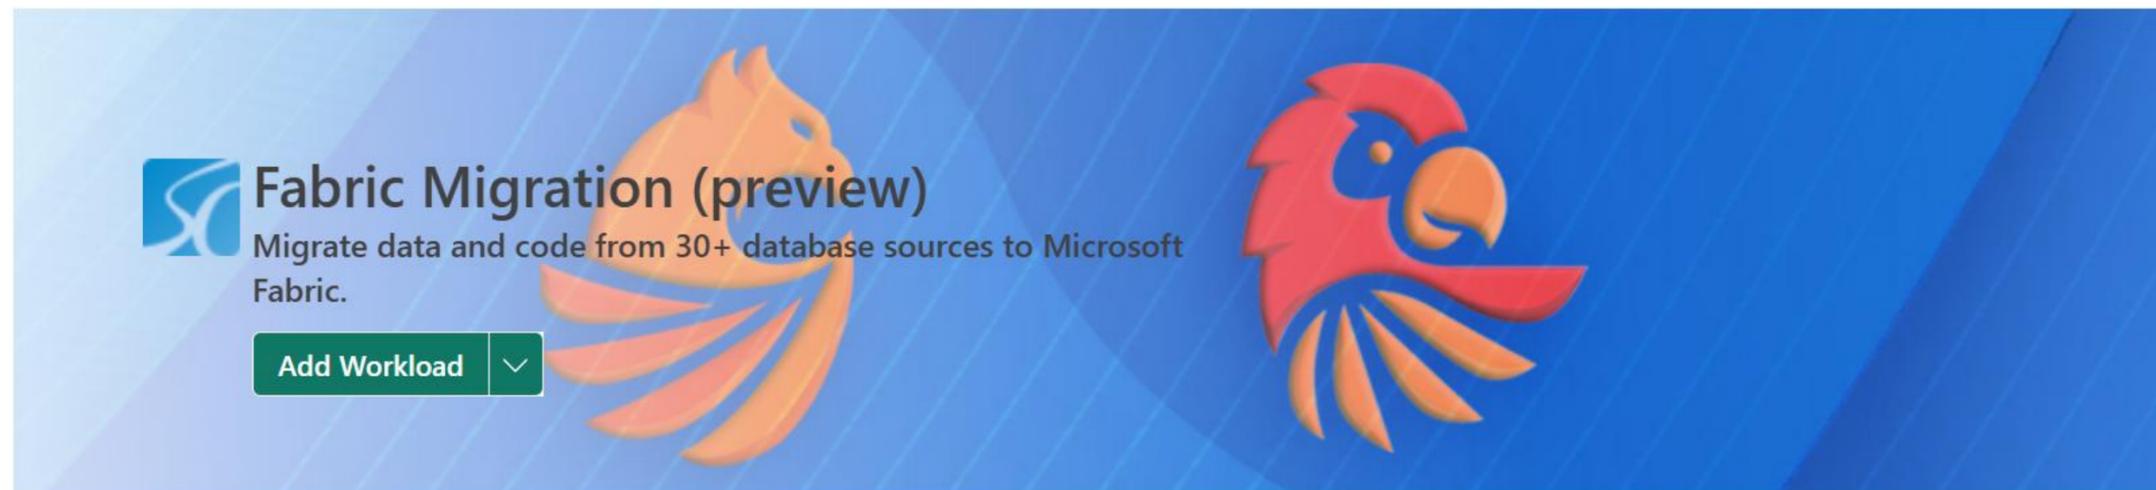
Kanerika Inc.'s migration accelerator helps move your Azure Data Factory (ADF) and Synapse pipelines into Fabric Data Factory (FDF)

- Assess & map ADF/Synapse pipelines
- Convert activities for Fabric optimization
- Configure integrations across lakehouse, warehouse & semantic models
- Validate, test, and optimize for peak performance

The screenshot shows the Microsoft Fabric marketplace interface. On the left is a navigation sidebar with icons for Home, Workspaces, Copilot, OneLake catalog, Monitor, Real-Time, Workloads, and My workspace. The main content area features a header for 'Azure to Fabric Migration (preview)' with the tagline 'The fastest way to migrate from Azure to Fabric with minimal downtime.' and an 'Add Workload' button. Below this are sections for 'About', 'Overview', and 'At a glance'. The 'About' section lists item types, compatibility with Lakehouse, and publisher support links. The 'Overview' section provides a description, publisher information (Kanerika Inc.), and license details. The 'At a glance' section includes a video player for 'Migrate Azure Pipelines to Microsoft Fabric | Smart Migration Accelerator | Data Migration Demo' with a 'Watch on YouTube' button and a 'Share' icon.



Third-party Fabric workloads



About

Item types



Compatible with



Publisher support

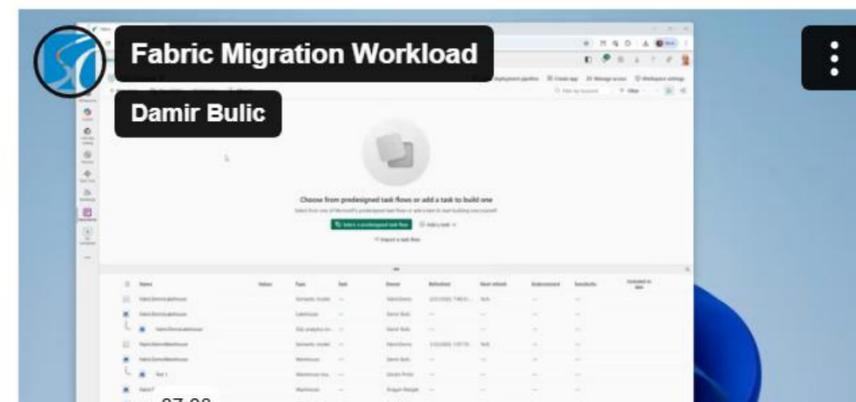
Overview

Description

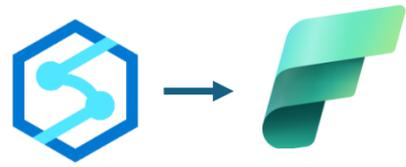
Publisher: Spectral Core

Spectral Core Fabric Migration is a native Microsoft Fabric workload that helps organizations migrate both data and database code to Fabric through one guided experience. It combines Omni Loader for high-performance data migration with SQL Tran for database code migration, including schemas, objects, SQL, and stored procedures, so teams can modernize existing platforms without

At a glance



- Announced this week
- Tool for migrating SQL Server and Synapse databases to Microsoft Fabric
- FREE plan supports:
 - SQL Server, Azure SQL Database, Azure SQL Managed Instance, and Synapse Analytics Dedicated SQL Pool.
 - Powerful testing framework.
- Premium plan adds:
 - Automated testing
 - Migration from other database sources



Microsoft Fabric Migration Resources



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



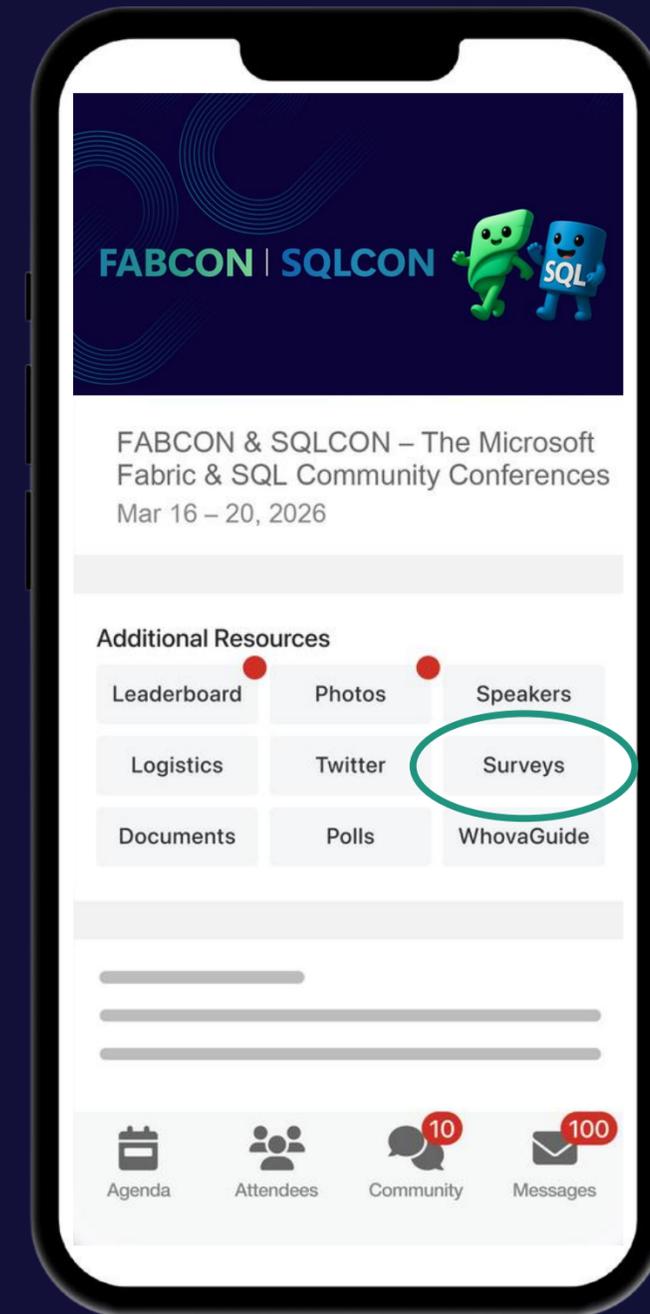
Influence our SQL roadmap and ensure it meets your real-life needs

<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>

