

#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





Data Warehouse Performance

1 TB TPC-H - 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 to Fabric Data Engineering from 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.)

AI-powered, multi-language notebooks with customizable environments for higher productivity


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

Upcoming Copilot-assisted capabilities for autonomous optimization and diagnostics



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 (Private Preview)



Data Engineering

Discovery and Assessment Tool

Migration Assistant (Private Preview)

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 predefined 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 predefined 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 VNET Data Gateway (if applicable) for bottlenecks
 - Use CETAS to export the data to ADLS Gen2 the 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 case insensitive collation
- Once a warehouse or SQL analytics endpoint is created, the collation cannot be changed

Review “State” of migrated objects

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

Review failed objects

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



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



Other Migration Considerations

aka.ms/fabricroadmap

T-SQL incompatibilities

- IDENTITY columns are BIGINT in Fabric Warehouse and you can't seed them
- Change GetDate() + n to DATEADD(), or CAST to datetime2, as datetime isn't supported
- HASHBYTES should use SHA2 instead of SHA1
 - String inputs should be varchar so you're hashing the exact UTF-8 string that will be stored
 - Regenerate these prior to migration and cast them to varchar first

ISOLATION LEVEL

- Fabric Warehouse supports SNAPSHOT isolation only; this blocks concurrent writes to a table
 - Stored procedures writing watermark values to a control table blocking each other
 - Either run these serially (slow) or change the code

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



What about Synapse Serverless SQL pools?

Migration Assistant only supports Dedicated SQL pools

- For Serverless Pools, assess and classify external tables/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 a **Table or Schema shortcut** under **/Tables**; the Lakehouse SQL endpoint will auto-discover these tables
- Shortcuts for **CSV/JSON/Parquet created** in **/Files** won't appear in the SQL endpoint
 - Ingest/transform them to Delta and write under **/Tables** (Dataflow Gen2 or Spark), then query via SQL

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


- Use the **SQL analytics endpoint** the same way:
 - Write tables as **Delta** under **/Tables** and query them.
- Migrating into OneLake reduces external dependencies, simplifies governance, and unlocks tight integrations like Direct Lake

Use a Warehouse when you need...

- **T-SQL data modification** (DML) and warehousing features beyond the endpoint's **read-only** scope
- Dedicated warehousing capabilities and management at Warehouse level
- **Tip:** start with *Lakehouse + SQL endpoint* (plus views) to replace Serverless-style querying and add a **Warehouse** only when you need writable T-SQL or Warehouse-specific features



Data Engineering migration





Items migration



Azure Synapse




Fabric Spark

Compute	Spark pools	➔ Starter and Custom Spark pools (memory optimized)
	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
	SJD pipeline activity	Fabric pipeline SJD activity
Data	ADLS Gen2 as main storage	• OneLake as main storage • Use shortcuts • ADLS Gen2
HMS	Internal	➔ • Use shortcuts > Delta tables only
	External (Azure SQL DB based)	• Lakehouse /Tables > Delta & non-Delta tables



Data Engineering Migration Assistant





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

λ

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

Q

Filter by keyword

Filter

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



How to Migrate: Additional Technical 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 data consumers (Power BI, apps, notebook, reports) to use Fabric endpoints

Monitoring & CI/CD

- Monitoring Hub, Spark Monitoring, Spark History Server
- CI/CD, incl ADO pipelines and GitHub actions integration

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.

Fabric Migration Assessment (preview)

Assess migration readiness

A built-in migration tool for upgrading to Microsoft Fabric will be available soon. In the meantime, you can use the [PowerShell script](#) to start your migration today. For additional guidance, check out our [migration overview](#) and [connector compatibility comparison](#) documentation.

Assessment Details

Review assessment for this Synapse workspace and inspect pipeline artifacts before migrating.

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.

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



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

DA

Data and Analytics

+

Overview

Boards

Repos

Files

Commits

Pushes

Branches

Tags

Pull requests

Advanced Security

Pipelines

Test Plans

Artifacts

Project settings

Synapse Prod

> credential

> dataflow

> dataset

> integrationRuntime

> linkedService

> managedVirtualNetwork

> notebook

> pipeline

> sparkConfiguration

> sqlscript

> trigger

publish_config.json

readme.md

master

Type to find a file or folder...

Files

Contents

History

Name ↑	Last change	Commits
credential	28 Jan	d0aa5
dataflow	28 Jan	d0aa5
dataset	28 Jan	d0aa5
integrationRuntime	Tuesday	f28a9
linkedService	10 Feb	3244f
managedVirtualNetwork	28 Jan	d0aa5
notebook	10 Feb	4b9b5
pipeline	6 Feb	2e8e0
sparkConfiguration	28 Jan	d0aa5
sqlscript	4 Feb	212ae
trigger	Monday	a005d
publish_config.json	28 Jan	f2796
readme.md	28 Jan	242b2

Initialized by Azure Synapse Analytics!

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



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](#)

Upload your ARM template to begin the migration process

[← 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 AI tools to identify
- GitHub Copilot with Fabric MCP worked well in the case of manual migrations

Use git integration and CI/CD

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



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

← Return to workloads

Azure to Fabric Migration (preview)

The fastest way to migrate from Azure to Fabric with minimal downtime.

[Add Workload](#)

About

Item types

[Azure to Fabric Migrati...](#)

Compatible with

[Lakehouse](#)

Publisher support

[Documentation](#) [Certification](#) [Help](#)

Overview

Description

Publisher: Kanerika Inc.

Kanerika's Azure to Fabric workload enables organizations to migrate Azure Data Factory and Azure Synapse Analytics pipelines to Microsoft Fabric in real-time.

License

View purchasing detail for Azure to Fabric Migration (preview) on the Azure Marketplace


[Azure marketplace offer page](#)

At a glance

Migrate Azure Pipelines to Microsoft Fabric | Smart Migration Acc... Data Migration Demo

Watch on [YouTube](#)

Azure to Microsoft Fabric



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
Uhova 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>

