



# Synapse to Fabric Migration Guide

Published: November 2025

## Table of Contents

Introduction .....	3
Why migrate from Synapse to Fabric? .....	3
Migration Guides.....	5
Synapse Data Warehouse Only .....	5
Synapse Spark Only .....	8
Hybrid with Synapse Spark and Synapse Data Warehouse .....	12
Unlock advanced capabilities in Microsoft Fabric.....	16
Continue learning with additional resources .....	17

## Introduction

This document outlines a migration strategy covering key workloads including Synapse dedicated SQL pools, Spark environments, and Synapse and Data Factory pipelines. It provides guidance on tooling and best practices to ensure a smooth transition with minimal disruption.

By following this strategy, organizations will be well positioned to unlock the full potential of Microsoft Fabric, enabling faster insights, streamlined operations, and a foundation for advanced analytics innovation.

## Why migrate from Synapse to Fabric?

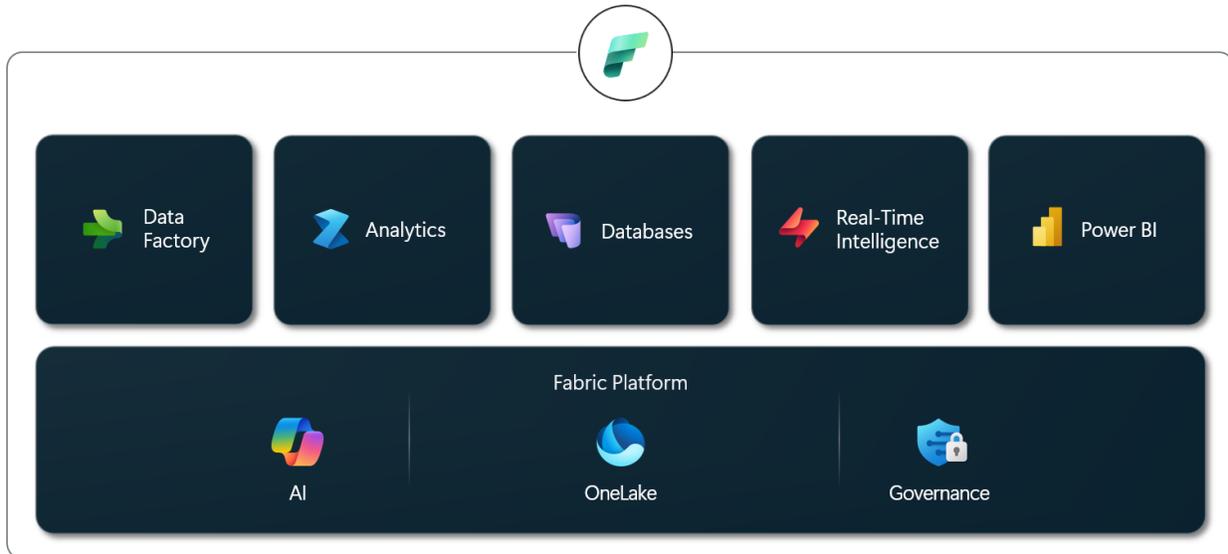
Microsoft Fabric introduced a new direction for analytics aimed at unification and simplicity. Fabric is positioned as the future of analytics, offering enhanced capabilities and performance for enterprises.

Microsoft Fabric is a unified, AI-powered data platform that bridges the gap between seasoned IT professionals and business users without formal coding expertise. For pros, it offers advanced capabilities like scalable data engineering, real-time analytics, and deep integration with Azure services, enabling complex workloads and custom solutions. Fabric empowers citizen developers through low-code/no-code experiences, intuitive orchestration tools, and seamless integration with familiar Microsoft 365 apps like Excel and Power BI. This dual approach removes traditional barriers, allowing all personas to collaborate on a single platform—accelerating innovation, improving governance, and democratizing access to enterprise-grade analytics.

Fabric supports conversational analytics and automated reasoning via tools like Fabric data agents and Copilot in Power BI. These features enable natural language interaction with diverse data types across OneLake and external sources.

# Microsoft Fabric

The unified data platform for AI transformation



- **Fabric Spark** is cheaper and faster in Fabric, thanks to features like the [Native Execution Engine](#) and [High Concurrency mode](#) - up to 3.5x faster than Synapse Spark and 3.7x faster than HDInsight for TPC-DS 1TB workloads.
- The redesigned **Fabric Data Warehouse** features an elastic, cloud-ready architecture, offering faster performance than Synapse Gen2 with less tuning and maintenance. Fabric data warehouse is up to 1.9x faster for TPC-H 1TB workloads and more than 5x faster for TPC-H 10TB workloads.
- **Fabric Data Factory** streamlines data integration with improved speed and ease, enhanced by Copilot, modern evaluators, parallel execution, copy job, and built-in activities for email and Teams.

Moving to Fabric is straightforward, thanks to its built-in features and tools designed to accelerate your migration.

- Data in ADLS Gen2 is instantly available in Fabric via OneLake shortcuts.
- Synapse pipelines can be moved to Fabric with PowerShell utilities.
- Spark workloads can be lifted and shifted to Fabric Spark using the fabric migration toolkit.
- Data warehouses are migrated automatically using our built-in, AI-assisted Migration Assistant.

Fabric offers greater speed, lower costs, enhanced security, and increased capabilities. Built for the era of AI, it's also built with your needs in mind.

## Migration Guides

Before embarking on any large-scale migration to Microsoft Fabric, it's essential to take a step back and assess your current architecture, workloads, and business goals.

Fabric introduces a unified, SaaS-based analytics platform that supports a wide range of data workloads, from structured data in warehouses to unstructured and semi-structured data in Lakehouses and Spark environments. While the platform offers significant performance, scalability, and integration benefits, successful migration requires thoughtful planning, especially when dealing with complex or high-volume systems.

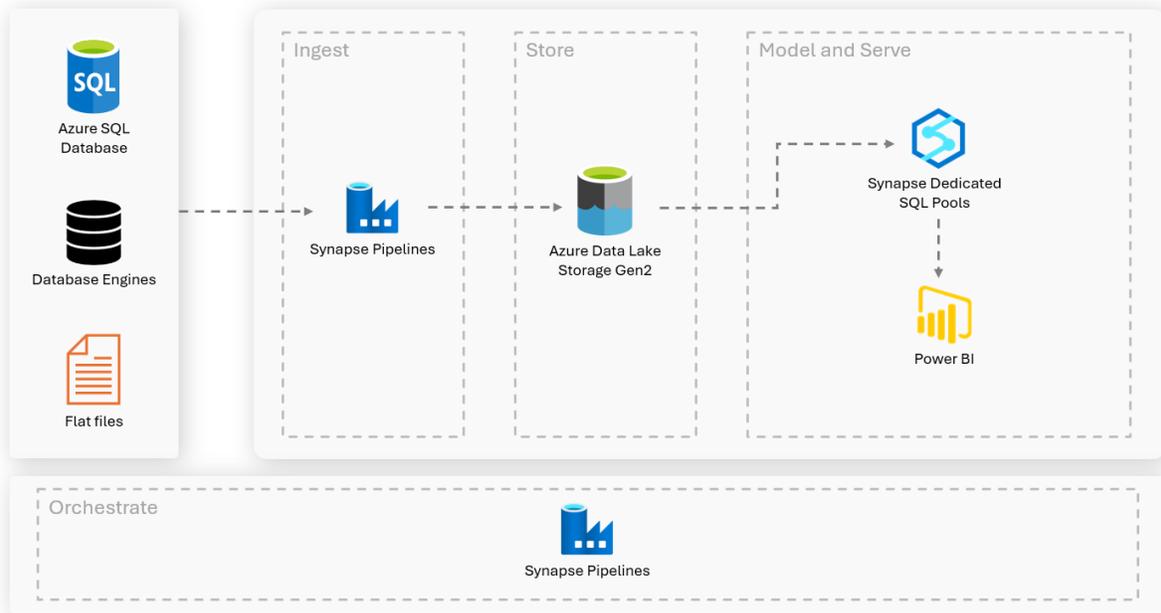
This section introduces a set of migration guides tailored to common architectural patterns, helping you navigate the journey from your current environment to Fabric with confidence.

If you're modernizing existing solutions or just lifting and shifting workloads, these guides will help you align your strategy with Fabric's capabilities and ensure a smooth transition.

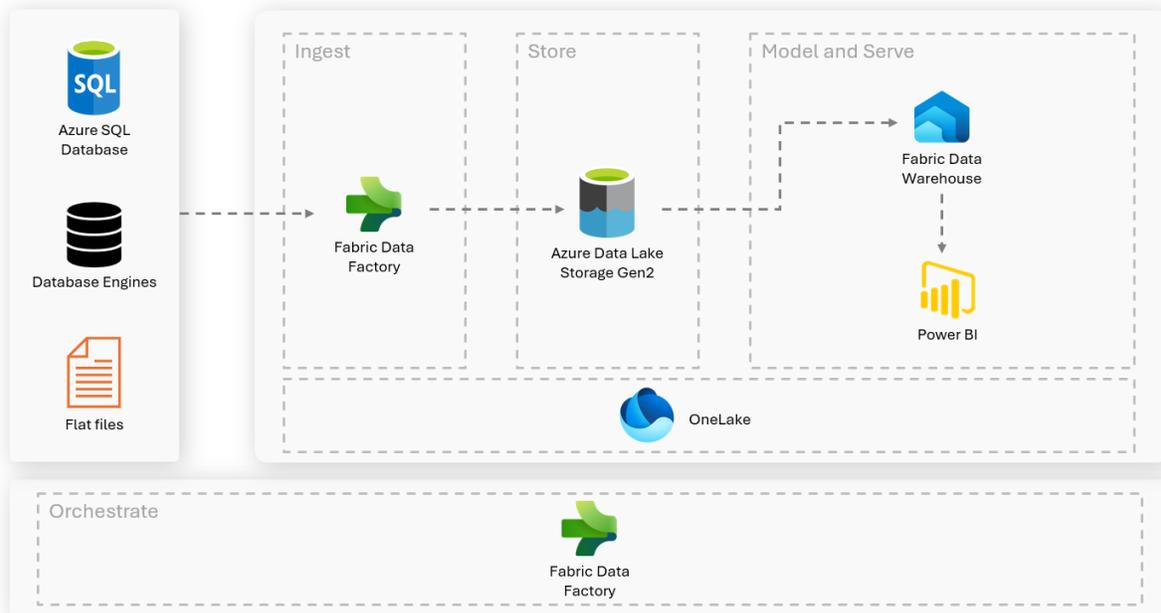
### Synapse Data Warehouse Only

Teams with strong T-SQL expertise or a background in building analytics solutions on relational databases - and working primarily with structured data - are well-suited for the data warehouse architecture. This approach often includes a data lake or a warehouse schema in the warehouse for long term raw data retention. Data prepping activities are largely handled through T-SQL logic in stored procedures while business users write ad hoc T-SQL queries against the warehouse.

This architecture typically moves data from source systems to an Azure Data Lake Storage Gen2 account(s) for raw data retention using Synapse pipelines, which is then loaded into a single dedicated SQL pool. Consumers include ad hoc queries, Power BI, and other analytical applications. All these processes are orchestrated by Synapse pipelines.



In a lift and shift to Fabric the dedicated SQL pool is replaced by the Fabric Data Warehouse and Synapse pipelines are replaced by Fabric Data Factory. The storage account could be moved to Fabric OneLake but it is advisable to perform this as a follow-on improvement to accelerate the migration process.



## Prerequisites

- Inventory objects by listing all tables, views, stored procedures, security roles, and other objects in your Synapse dedicated pool.

- Inventory all upstream data sources and downstream data consumers that will need to have connections rerouted.

## Key Considerations

- Identify [incompatible T-SQL](#) features, [DMVs](#), or [security](#) configurations that need refactoring for Fabric.
- Review the migration considerations, strategies, and approaches for moving [Synapse pipeline to Fabric](#).
- Start with a smaller workload (1-10 TB) to validate migration and create a repeatable template for future migrations.

## What components should be migrated?

PaaS Service	Fabric Item
Dedicated SQL Pools Gen2	Data Warehouse
Synapse pipelines	Data Factory
Azure Data Lake Storage Gen2	Lakehouse (Files section)

## Migration steps

The recommended steps for migration are as follows:

- 1. Environment Setup:** Determine the capacity size needed using the [SKU Estimator](#) and provision a Fabric capacity or use the 60-day [free trial](#), create a workspace, setup private networking as required, and ensure at least *member* permission has been granted to all individuals involved in the initial migration effort.
- 2. Migrate the dedicated SQL pool:** Extract a dacpac of the Synapse dedicated SQL pool Gen2 which will include the definition of all the objects and security but does not contain any user data. Then use the built-in [Fabric Migration Assistant](#) which will convert the schema and provide AI powered fixes for incompatible code. Remediate any outstanding code issues.
- 3. Migrate the data:** Using either Data Migration Assistant or a set of manually built one-time use Fabric Data Factory pipelines or a copy job, copy the data from the Synapse dedicated SQL pool Gen2 to the newly created Fabric Data Warehouse. Validate data row counts and integrity between the two environments after the initial copy is completed.
- 4. Migrate the Synapse pipelines:** Using the [Fabric Pipeline Upgrade PowerShell module](#), convert Data Factory pipelines to the corresponding Fabric pipeline JSON for the supported capabilities or use the [Data Factory Migration Assistant](#) to automate the migration of pipelines then fill in the gaps by manually moving components, rearchitecting as necessary.

5. **Validate security:** The security structure differs slightly between Synapse and Fabric. Review the Synapse workspace level or logical Azure SQL Server permissions as appropriate to ensure the proper level of access is granted to users which will get workspace level permissions. Additionally, plan to replace any SQL authentication with Entra ID. For a comprehensive look at securing Fabric review the [Microsoft Fabric Security](#) documentation.
6. **Reroute connections:** Using the inventory of data consumers, reroute connections from the Synapse dedicated SQL pool to the Fabric Data warehouse.
7. **Parallel run:** Optionally, run the two environments in parallel for a period to ensure all ETL processes are functioning properly, and data consumers are aware of the connection changes.

## Optimize and Extend

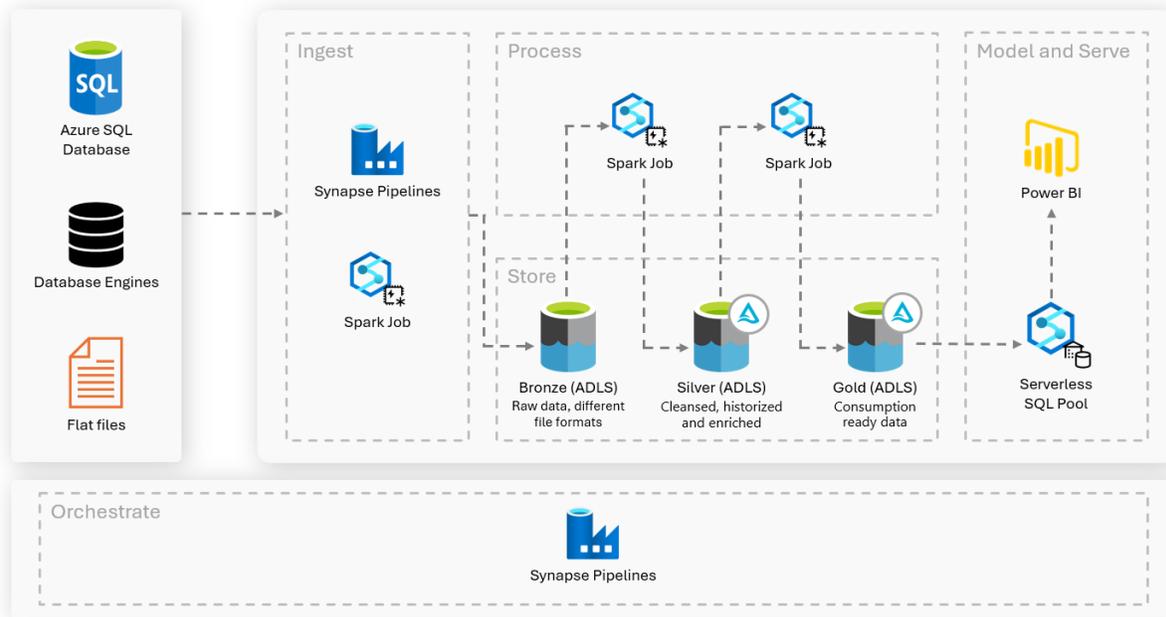
For this specific architecture consider:

- When migrating the data, using the Migration Assistant's built-in capabilities is simple and easy for smaller databases (~1TB or less). If you require higher throughput, consider the following alternative option:
  - Extract the data via [CREATE EXTERNAL TABLE AS SELECT](#) (CETAS) from the dedicated SQL pool to Azure Data Lake Storage Gen2.
  - Use either [Fabric Data Factory](#) or the [T-SQL COPY command](#) to ingest the data into Fabric Warehouse.

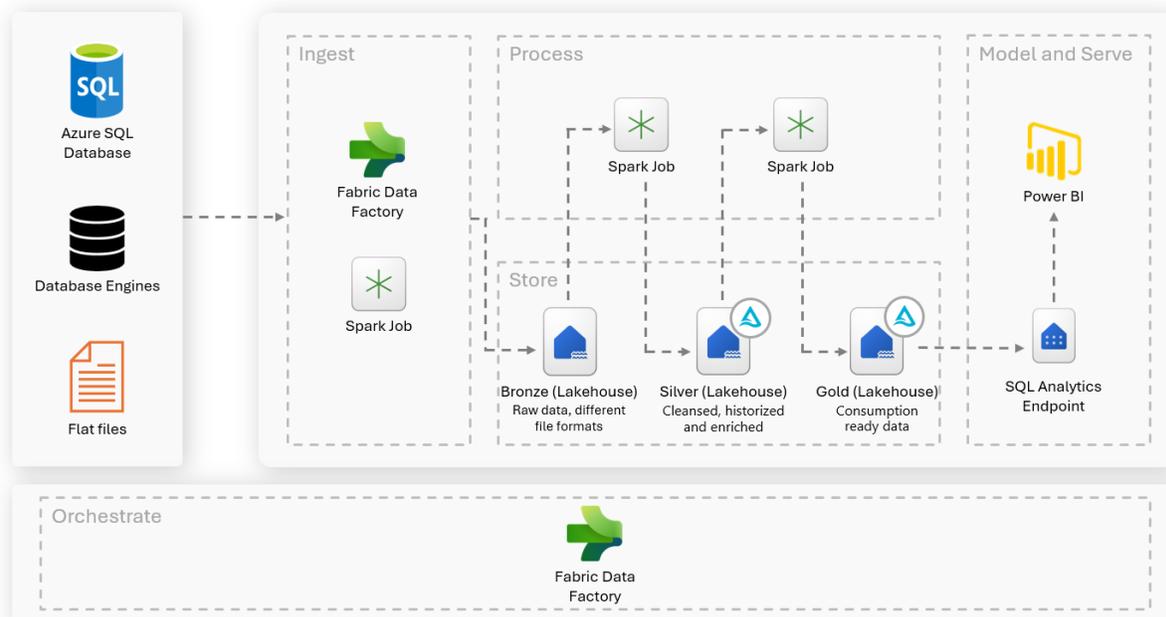
## Synapse Spark Only

Teams proficient in Apache Spark (Scala, Python, Spark SQL) with complex data transformation requirements or with heavy semi-structured and/or unstructured data requirements would be best fit to leverage the Lakehouse Item in Fabric for all data zones. The SQL Analytics Endpoint provides interactive and scalable querying capabilities to empower all Lakehouse data consumers.

The Lakehouse architecture typically moves data from source systems to an Azure Data Lake Storage Gen2 account(s) for raw data retention using a combination of Synapse pipelines and Spark jobs. Raw data is then cleaned and standardized using Spark jobs before landing in the silver layer in the Delta format. Finally, data products are stored in the gold layer in Delta format and served to end users through the Synapse serverless SQL pools, Power BI, and other analytical applications. All these processes are orchestrated by Synapse pipelines.



In a lift and shift to Fabric the Azure Data Lake Storage Gen2 is replaced by the Fabric Lakehouse, the Synapse serverless SQL pool is replaced by the lakehouse's built-in SQL analytics endpoint, Synapse Spark jobs are replaced by Fabric Spark jobs, and Synapse pipelines are replaced by Fabric Data Factory.



## Prerequisites

- Inventory all upstream data sources and downstream data consumers that will need to have connections rerouted.

- Inventory Spark pools, jobs, and notebooks being used in the Synapse environment.

## Key Considerations

- Review [feature comparison](#) and [roadmap](#) for Fabric Data Engineering vs. Azure Synapse Spark.
- Review the migration considerations, strategies, and approaches for moving [Synapse pipeline to Fabric](#).

## What components should be migrated?

PaaS Service	Fabric Item
Synapse pipelines	Data Factory
Azure Data Lake Storage Gen2	Lakehouse
Spark job definition	Spark job definition
Spark notebook	Spark notebook
Serverless SQL Pool	Lakehouse (SQL analytics endpoint)
Spark pool	Spark pool
Spark configurations	Environment
Libraries	Environment

## Migration steps

The recommended steps for migration are as follows:

- 1. Environment Setup:** Determine the capacity size needed using the [SKU Estimator](#) and provision a Fabric capacity or use the 60-day [free trial](#), create a workspace, setup private networking as required, and ensure at least *member* permission has been granted to all individuals involved in the initial migration effort.
- 2. Create the Fabric Lakehouse(s):** Many data engineering items in Fabric require a lakehouse context to run. Create a lakehouse per zone in the architecture or a single lakehouse depending on your use environment's configuration.
- 3. Migrate the data:** With the lakehouse(s) in place, copy the data from the Azure Data Lake Storage Gen2 account(s) to the corresponding Fabric lakehouse(s). Files (non-Delta) can be moved using [AzCopy](#), Fabric pipelines, or a Fabric Spark notebook. Move Delta tables using Fabric Spark notebooks that read the table and write it back to the lakehouse Tables section. An alternative to physically moving the data is to use shortcuts to access the data directly in Azure Data Lake Storage Gen2.
- 4. Migrate Spark components:** Using the [migration toolkit](#), migrate each component using the provided scripts where available. Migrate items in the order listed below. You can also use Fabric REST API to move pools, libraries and configurations.

Item	Can migrate manually	Can script/API migration
Pools	✓	✓
Configuration	✓	✓
Libraries	✓	✓
Notebooks	✓	✓
Spark job definition	✓	✓
HMS metastore	✓	✓

5. **Migrate the Synapse serverless pool:** Many, if not all, tables will already be available in the SQL analytics endpoint automatically. For tables that are not automatically discovered, and which may be stored elsewhere, can be recreated manually by creating shortcuts in the lakehouse Tables section or using the [OPENROWSET](#) T-SQL function in a view. Views should be checked and manually modified as necessary.
6. **Migrate the Synapse pipelines:** Using the [Fabric Pipeline Upgrade PowerShell module](#), convert Data Factory pipelines to the corresponding Fabric pipeline JSON for the supported capabilities or use the [Data Factory Migration Assistant](#) to automate the migration of pipelines then fill in the gaps by manually moving components, rearchitecting as necessary.
7. **Validate security:** The security structure differs slightly between Synapse and Fabric. Review the Synapse workspace level as appropriate to ensure the proper level of access is granted to users which will get workspace level permissions. Additionally, plan to replace any SQL authentication with Entra ID. For a comprehensive look at securing Fabric review the [Microsoft Fabric Security](#) documentation.
8. **Reroute connections:** Using the inventory of data consumers, reroute connections from the Synapse serverless SQL endpoint to the Fabric lakehouse's SQL analytics endpoint.
9. **Parallel run:** Optionally, run the two environments in parallel for a period to ensure all ETL processes are functioning properly, and data consumers are aware of the connection changes.

## Optimize and Extend

For this specific architecture consider:

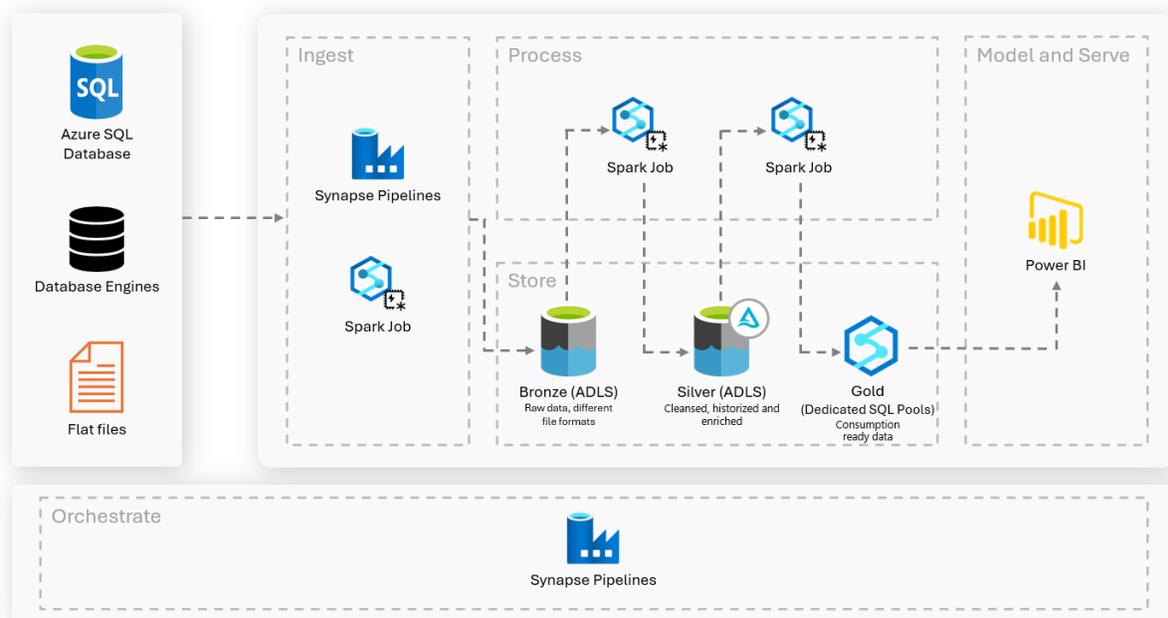
- Explore the [Native Execution Engine](#) (currently in preview), which accelerates job execution through optimized task planning and improved I/O pathways.

- Use the [Autotune feature](#) to automatically adjust Spark configurations for specific workloads by leveraging historical job runs.

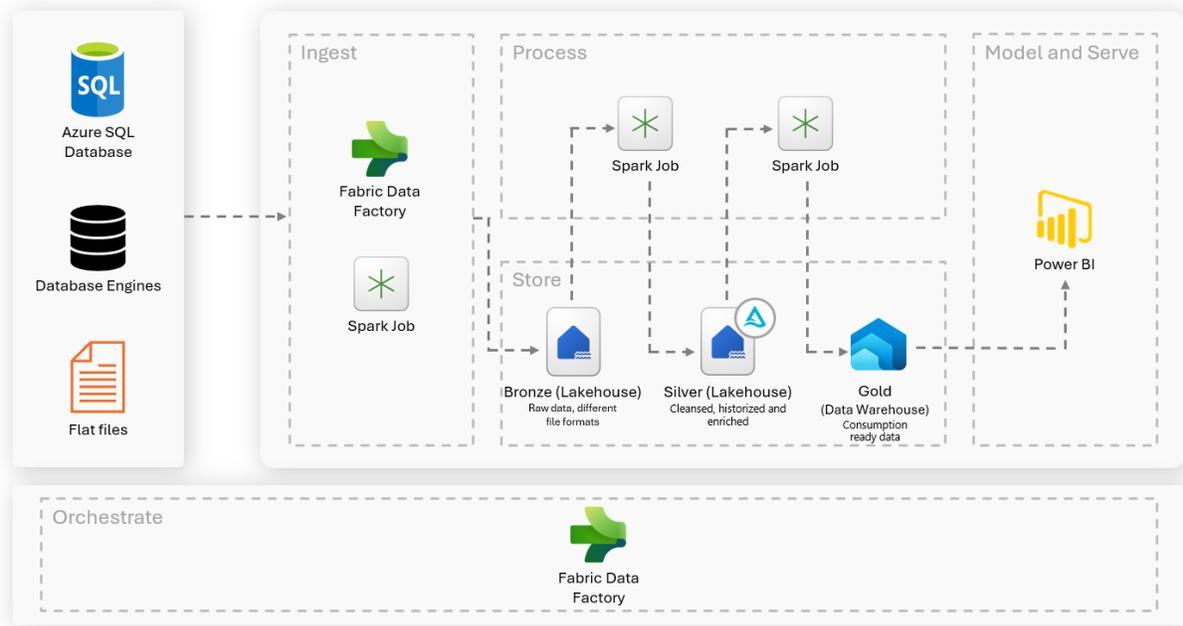
## Hybrid with Synapse Spark and Synapse Data Warehouse

Teams with mixed skillsets, or multiple teams collaborating on single unified architecture may be best suited for a hybrid architecture with both Spark and the Fabric Warehouse being leveraged based on responsibility boundaries that may exist between creator personas.

This architecture typically moves data from source systems to an Azure Data Lake Storage Gen2 account(s) for raw data retention using a combination of Synapse pipelines and Spark jobs. Raw data is then cleaned and standardized using Spark jobs before landing in the silver layer in the Delta format. Finally, data products are stored in the gold layer in Delta format and served to end users through the Synapse serverless SQL pools, Power BI, and other analytical applications. All these processes are orchestrated by Synapse pipelines.



In a lift and shift to Fabric the Azure Data Lake Storage Gen2 is replaced by the Fabric Lakehouse, the Synapse dedicated SQL pool Gen2 is replaced by the data warehouse, Synapse Spark jobs are replaced by Fabric Spark jobs, and Synapse pipelines are replaced by Fabric Data Factory.



## Prerequisites

- Inventory objects by listing all tables, views, stored procedures, security roles, and other objects in your Synapse dedicated pool.
- Inventory all upstream data sources and downstream data consumers that will need to have connections rerouted.
- Inventory Spark pools, jobs, and notebooks being used in the Synapse environment.

## Key Considerations

- Identify [incompatible T-SQL](#) features, [DMVs](#), or [security](#) configurations that need refactoring for Fabric.
- Review [feature comparison](#) and [roadmap](#) for Fabric Data Engineering vs. Azure Synapse Spark.
- Review the migration considerations, strategies, and approaches for moving [Synapse pipelines to Fabric](#).

## What components should be migrated?

PaaS Service	Fabric Item
Synapse pipelines	Data Factory
Azure Data Lake Storage Gen2	Lakehouse
Spark job definition	Spark job definition
Spark notebook	Spark notebook
Dedicate SQL pool Gen2	Data warehouse
Spark pool	Spark pool

Spark configurations	Environment
Libraries	Environment

## Migration steps

The recommended steps for migration are as follows:

- 1. Environment Setup:** Determine the capacity size needed using the [SKU Estimator](#) and provision a Fabric capacity or use the 60-day [free trial](#), create a workspace, setup private networking as required, and ensure at least *member* permission has been granted to all individuals involved in the initial migration effort.
- 2. Create the Fabric Lakehouse(s):** Many data engineering items in Fabric require a lakehouse context to run. Create a lakehouse per zone in the architecture or a single lakehouse depending on your use environment's configuration.
- 3. Migrate the lakehouse data:** With the lakehouse(s) in place, copy the data from the Azure Data Lake Storage Gen2 account(s) to the corresponding Fabric lakehouse(s). Files (non-Delta) can be moved using [AzCopy](#), Fabric pipelines, or a Fabric Spark notebook. Move Delta tables using Fabric Spark notebooks that read the table and write it back to the lakehouse Tables section. An alternative to physically moving the data is to use shortcuts to access the data directly in Azure Data Lake Storage Gen2.
- 4. Migrate Spark components:** Using the [migration toolkit](#), migrate each component using the provided scripts where available. Migrate items in the order listed below. You can also use Fabric REST API to move pools, libraries and configurations.

Item	Can migrate manually	Can script/API migration
Pools	✓	✓
Configuration	✓	✓
Libraries	✓	✓
Notebooks	✓	✓
Spark job definition	✓	✓
HMS metastore	✓	✓

- 5. Migrate the dedicated SQL pool:** Extract a dacpac of the Synapse dedicated SQL pool Gen2 which will include the definition of all the objects and security but does not contain any user data. Then use the built-in [Fabric Migration Assistant](#) which will convert the schema and provide AI powered fixes for incompatible code. Remediate any outstanding code issues.
- 6. Migrate the dedicated SQL pool data:** Using either Data Migration Assistant or a set of manually built one-time use Fabric Data Factory pipelines or a copy job, copy the data from the Synapse dedicated SQL pool Gen2 to the newly created Fabric

Data Warehouse. Validate data row counts and integrity between the two environments after the initial copy is completed.

7. **Migrate the Synapse pipelines:** Using the [Fabric Pipeline Upgrade PowerShell module](#), convert Data Factory pipelines to the corresponding Fabric pipeline JSON for the supported capabilities or use the [Data Factory Migration Assistant](#) to automate the migration of pipelines then fill in the gaps by manually moving components, rearchitecting as necessary.
8. **Validate security:** The security structure differs slightly between Synapse and Fabric. Review the Synapse workspace level or logical Azure SQL Server permissions as appropriate to ensure the proper level of access is granted to users which will get workspace level permissions. Additionally, plan to replace any SQL authentication with Entra ID. For a comprehensive look at securing Fabric review the [Microsoft Fabric Security](#) documentation.
9. **Reroute connections:** Using the inventory of data consumers, reroute connections from the Synapse dedicated SQL pool to the Fabric Data warehouse.
10. **Parallel run:** Optionally, run the two environments in parallel for a period to ensure all ETL processes are functioning properly, and data consumers are aware of the connection changes.

## Optimize and Extend

For this specific architecture consider:

- Explore the Native Execution Engine (currently in preview), which accelerates job execution through optimized task planning and improved I/O pathways.
- Use the Autotune feature to automatically adjust Spark configurations for specific workloads by leverages historical job runs.
- When migrating the data, using the Migration Assistant's built-in capabilities is simple and easy for smaller databases (~1TB or less). If you require higher throughput, consider the following alternative option:
  - Extract the data via [CREATE EXTERNAL TABLE AS SELECT](#) (CETAS) from the dedicated SQL pool to Azure Data Lake Storage Gen2.
  - Use either [Fabric Data Factory](#) or the [T-SQL COPY command](#) to ingest the data into Fabric Warehouse.

# Unlock advanced capabilities in Microsoft Fabric

Once your migration from Azure Synapse Analytics to Microsoft Fabric is complete, it's time to go beyond the basics and take full advantage of Fabric's integrated ecosystem. These features will help you modernize your analytics workflows and maximize value:

## **Shortcuts**

Create virtual links to external data without copying it, enabling seamless access across multiple sources in OneLake. This reduces storage duplication and accelerates data integration.

## **Mirroring**

Keep your data synchronized in near real-time between external systems and Fabric. Mirroring supports SQL Server, Azure SQL Database, and other sources, providing a consistent view for analytics and reporting.

## **Data Agents**

Use generative AI to let users interact with organizational data in a conversational Q&A experience, pulling insights from sources like lakehouses, warehouses, Power BI semantic models, and KQL databases without writing complex queries.

## **Power BI Direct Lake**

Deliver lightning-fast insights by reading data directly from OneLake without needing to import data or run queries against the source data. Direct Lake mode combines the flexibility of self-service BI with enterprise-grade performance.

## **SQL Database in Fabric**

Enjoy a familiar T-SQL experience within Fabric for building relational solutions. SQL Database integrates seamlessly with other Fabric services, making it easy to combine structured and unstructured data in one environment.

## **Dataflows Gen2**

Low-code data preparation using the Power Query experience and AI-powered Copilot to connect, transform, and load data into destinations like Lakehouse, Warehouse, or Azure SQL with streamlined authoring and automation.

## Continue learning with additional resources

Learn about available migration tooling:

- [Signup for the Synapse Spark to Fabric Spark Migration Tool Private Preview](#)
- [Fabric Data Warehouse Migration Assistant](#)
- [Migrate pipelines using PowerShell](#)

Useful resources for working with and learning Microsoft Fabric:

- [What is Microsoft Fabric? | Microsoft Learn](#)
- [Microsoft Fabric Roadmap](#)
- [Data Engineering in Microsoft Fabric documentation | Microsoft Learn](#)
- [Data Engineering Tutorial | Microsoft Learn](#)
- [Migrating from Azure Synapse Spark to Fabric | Microsoft Learn](#)
- [Comparison between Fabric and Azure Synapse Spark | Microsoft Learn](#)
- [Data Factory in Microsoft Fabric documentation | Microsoft Learn](#)
- [Data Factory Tutorial | Microsoft Learn](#)
- [Data Warehouse in Microsoft Fabric documentation | Microsoft Learn](#)
- [Data Warehouse Tutorial | Microsoft Learn](#)
- [Fabric Toolbox | GitHub](#)
- [Fabric SKU Estimator](#)