SPYGLASSMTG

A Women Owned/Women Led Company

Spyglass Presents:
Fabric Solutions in 30 Days (FAi30)

Microsoft
Solutions Partner

Infrastructure Azure
Data & AI Azure
Digital & App Innovation Azure
Modern Work
Security

Spyglass's Fabric Solution in 30 Days (FAi30)

- **What is FAi30?**
  - Fabric Solution in FAi30 is designed to help you create a Fabric Solution in Azure quickly and efficiently.
  - With defined engineering approaches to support the Fabric "Lakehouse" or "Warehouse" sewing data together is quick and easy
  - With dozens of prebuilt templates, adoption guidance, governance standards and development standards; your solution will be up and running before you know it

- **FAi30 will allow your business to:**
  - Integrate multiple data sources to create a Lakehouse and/or Warehouse that supports your business initiatives
  - Quickly integrate intelligent apps, backed by Azure's prebuilt AI capabilities into Fabric
  - Gain a complete understanding of Fabric, its capabilities and best practices
  - Establish an MVP governance and adoption framework to enable self-service analytics

SPYGLASS MTG

Spyglass FAi30 Approach

1. **Strategy & Plan: Governance, Analytics and Architecture Workshops**
   - Identify the who, what, when, and why including inventory review.
   - Envision the role of Fabric solutions in your organization
   - Identification of Advanced Analytical Scenarios
   - Analyze the customer data estate, infrastructure, applications, workloads, and data scenarios to scope the transition requirement
   - Analyze the compliance, security, and governance implications of Fabric

2. **Ready & Govern:** Implement Fabric Governance & Security
   - Implement the MVP Fabric framework for data governance
   - Implement the MVP security framework for workspace and data asset management
   - Establish connectivity for the data solution

3. **Design & Innovate**: Model the data (decide how the data you need and create will be organized, accessed, and stored) and sketch out the app or report screens.
   - Design your Fabric data solution to host your data
   - Design new Business Intelligence insights
   - Design ongoing Fabric and data governance
   - Review additional scenarios and develop roadmap to an optimized environment based on business goals
   - Build and Implement the Fabric solution

4. **Transition:** Get your solution into the hands of users, get feedback, and decide what to change or add.
   - Training and hand-off throughout the project lifecycle

SPYGLASSMTG

FAi30 Toolkit

- Governance First Adoption Framework
- Solution Design Template
- Planning and Delivery Checklists
- Documented Development Practices
- Documented Architecture Design Patterns
- Capacity Assessment and Planning Tools
- Security and Compliance Checklists
- Performance Management Checklist
- Experienced Deployment Engineers
- Fabric Solution Delivery Templates

SPYGLASSMTG

Spyglass FSi30 Benefits

- Spyglass FAi30 provides the following benefits:
  - **Governance –** establish a governance framework the right way; before it gets out of hand!

  - **Time to Market** – faster and more accurate insights and predictions 2-10x faster development ROI

  - **Flexibility** – Works with your current data sources to get the most from existing investments and grows with your business

  - **Repeatability** – Spyglass's proven methodology and templates provides repeatable process and accelerators to speed the development of your solution

SPYGLASSMTG

Next Steps

Scoping Session to understand client's data landscape
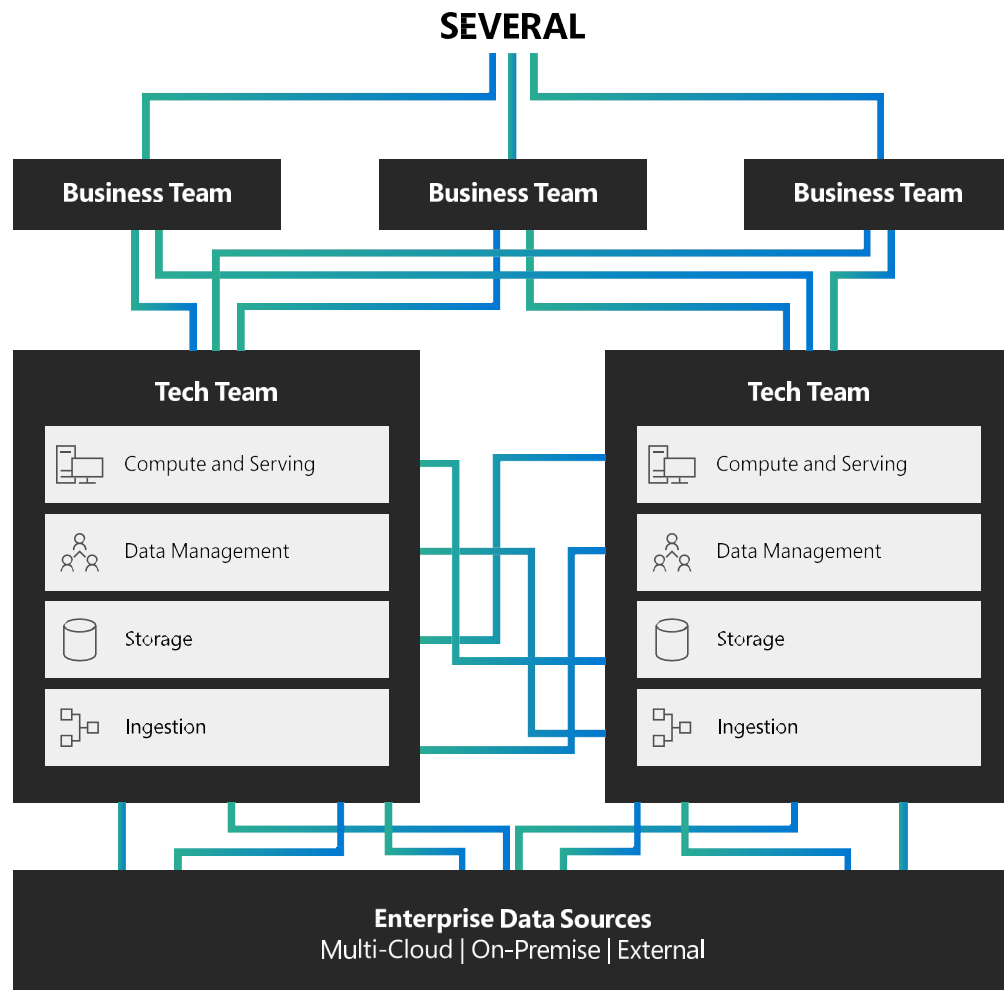
Process paperwork

Start building solution!

SPYGLASS MTG

# Thank You

While data grows 400%, less than 30% gets analyzed

44ZB

175ZB

SPYGL\SSMTG

An organically evolved data estate?

SEVERAL

Business Team | Business Team | Business Team

Tech Team
- Compute and Serving
- Data Management
- Storage
- Ingestion

Tech Team
- Compute and Serving
- Data Management
- Storage
- Ingestion

Enterprise Data Sources
Multi-Cloud | On-Premise | External

The most common challenge in enterprise data estates is that data has organically evolved.

Companies have data estates with a ton of data and infrastructure redundancy. There are often multiple teams operating silos of data that are not truly connected. The platforms required to transform data into actionable insights are inherently very technical and fragmented. Usually, teams of engineers and developers are needed to deploy and manage these platforms. Organizations have many siloed cloud or on-premises data sources from different vendors in different formats that hold critical information.   It is very difficult to find deep and accurate insights without a single source of truth. Stitching together the unique analytics tools organizations need is complicated. Costs associated with procuring and managing these capabilities can be exorbitant. And there is a significant risk associated with the lack of governance.

- **Every analytics project has many subsystems**
- **Every subsystem needs a different class of product**
- **Products often come from multiple vendors**
- **Integration across products is complex, fragile and expensive**

SPYGLASSMTG

Analytics & AI is the #1 investment for business leaders, however they struggle to maximize ROI

?

80%

55%

10

SPYGLASS MTG

More analytics solutions lead to more silos

Data
Structured
Unstructured
Streaming
Big data
Cloud
On-premises
IoT/edge

Technologies
Map Reduce
Open source projects
Data mart
Data warehouse
Data lake
Spreadsheets
RDMS
Data visualization tools
ML models
AI services

Skills
SQL
Python
Java
R
Industry Schemas
Data modeling
Data cleaning
Data integrations

SPYGLASSMTG

Analytics should seamlessly be part of the way users work—rather than a factor for creating more siloes

SPYGLASS MTG
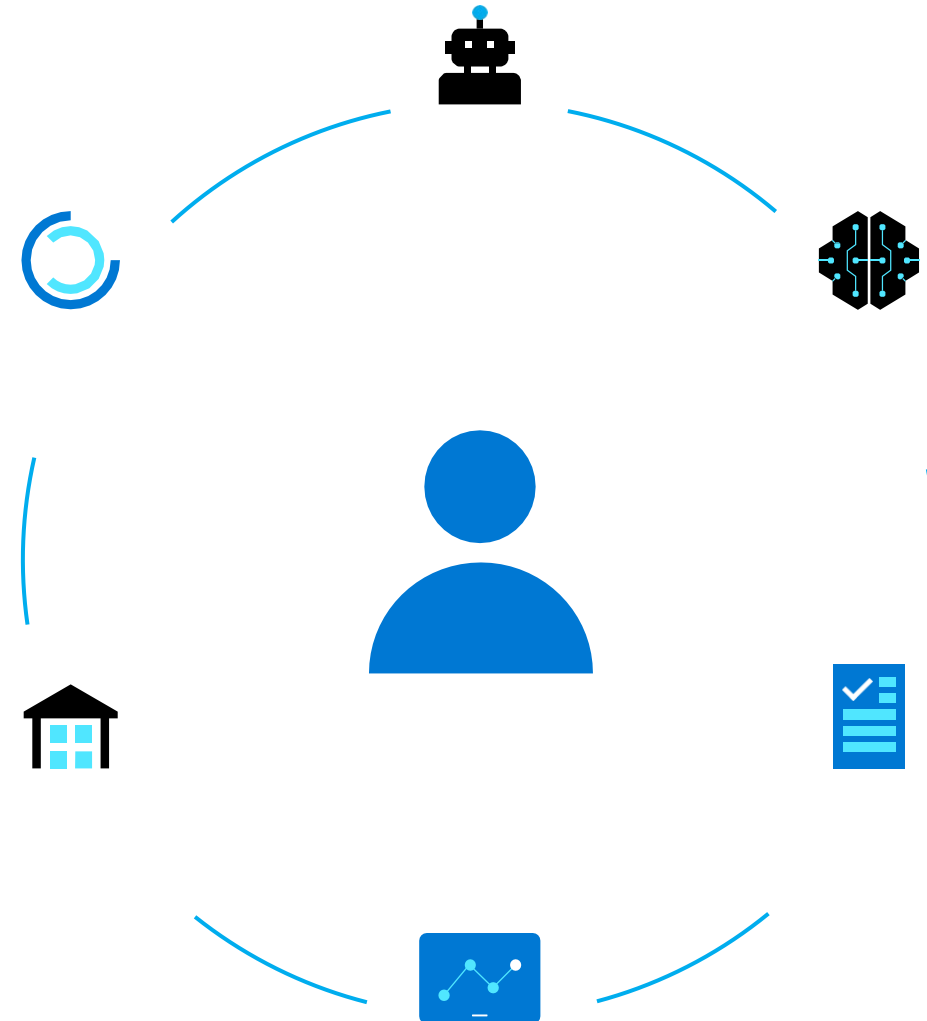
Microsoft Fabric

**Analytics Platform as a Service
with unmatched time to insight**

Data Mesh, Data Fabric, Data Hub

There are three data estate architectures and concepts that organizations are applying to the modernization of their data estate. The data mesh, data fabric, and data hub.

Data Mesh

A series of domains assigned to individual LOBs that enables access to the data they need with maximum autonomy by upholding the four principles of a data mesh.

Data Fabric

A system for automating data management tasks, such as unifying and cleaning disparate sources as well as authorizing data access, that helps a business make the most of its existing data sources without needing to migrate them.

Data Hub

An open and governed lakehouse platform that provides the storage foundation for serving data to multiple domains efficiently, so domain users can easily consume it.

SPYGLASSMTG

Data Landing Zones

Data landing zone architecture illustrates the layers, their resource groups, and services each resource group contains. The architecture also provides an overview of all groups and roles associated with your data landing zone, plus the extent of their access to your control and data planes.

Data Management

The data management landing zone is a management function and is central to cloud-scale analytics. It's responsible for the governance of your analytics platform.

Data Integration

Each data landing zone has several layers that enable agility for the service data integrations and data products it contains. You can deploy a new data landing zone with a standard set of services that let the data landing zone begin ingesting and analyzing data.

Data Products

Data products are data served as product and computed, saved, and served by polyglot persistence services, which can be required by certain use cases. The process of creating and serving a data product can require services and technologies that aren't included in the data landing zone core services.

SPYGLASSMTG

Next Gen: Microsoft Fabric
A unified analytics solution for the era of AI

Enabling the art of the possible, exceeding customer expectations, and delivering transformative business value just got easier with the introduction of Microsoft Fabric.

Microsoft Fabric brings together **the best parts of data mesh, data fabric, and data hub** to provide **a one-stop-shop for data integration, data engineering, real-time analytics, data science, and business intelligence** needs without compromising the privacy and security of your data. By joining top companies using Microsoft Fabric, you will enable teams to experience an end-to-end, highly integrated, single solution that is easy to understand, onboard, create, and operate.

With this new solution, you will establish a **unified source of truth** by bringing together all analytics workloads in a lake-first foundation. Teams will also be able to reduce the time and effort it takes to uncover impactful insights through **democratized access** to data. And this can be done confidently through a secure, **governed solution**.

Data Integration
This solution offers comprehensive data migration and integration to enable a lake-first pattern. Azure Data Factory connectors enable data integration, while Azure Synapse Link connectors enable "no code" and "always synchronized" data integration for operational databases.

OneLake
All data is ingested into a data lake on Azure Data Storage Gen––a cost- and performance-optimized data lake storage service—for the most demanding business intelligence, machine learning, and artificial intelligence workloads.

Analytics
Data scientists can bring their preferred compute frameworks, languages, runtimes, and tools into the data lakehouse and further enhance the data through feature engineering and statistical techniques.

Business Intelligence
Best-in-class integrated solutions to responsibly democratize business intelligence with self-serve tools and experiences for data analysts and data citizens.

Governance
Microsoft Purview then provides a single pane governance solution to help effectively scan and manage your data estate—even as it grows and scales.
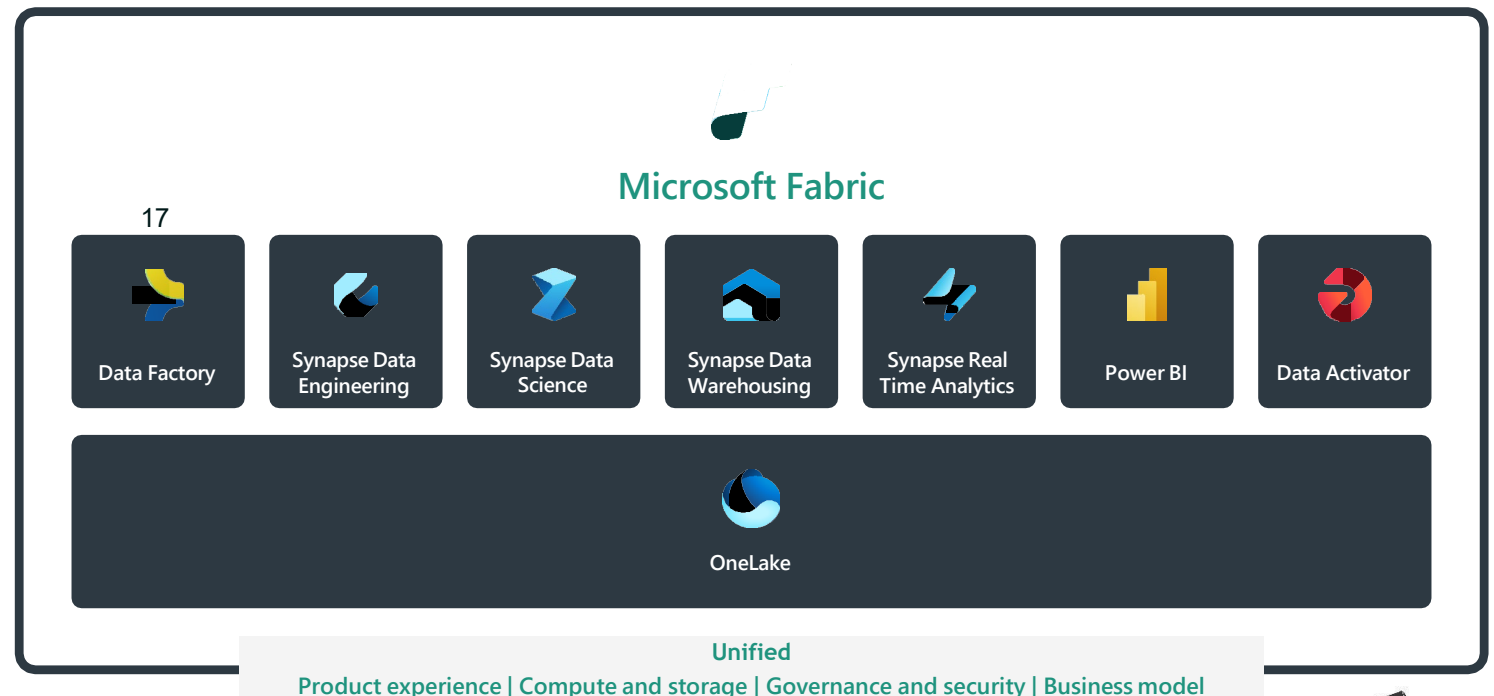
A unified SaaS-based solution that stores all organizational data where analytics workloads operate

Microsoft Fabric combines Data Factory, Synapse Analytics, Data Explorer, and Power BI into a single, unified experience, on the cloud. The open and governed data lakehouse foundation is a cost- effective and performance-optimized fabric for business intelligence, machine learning, and AI workloads at any scale.
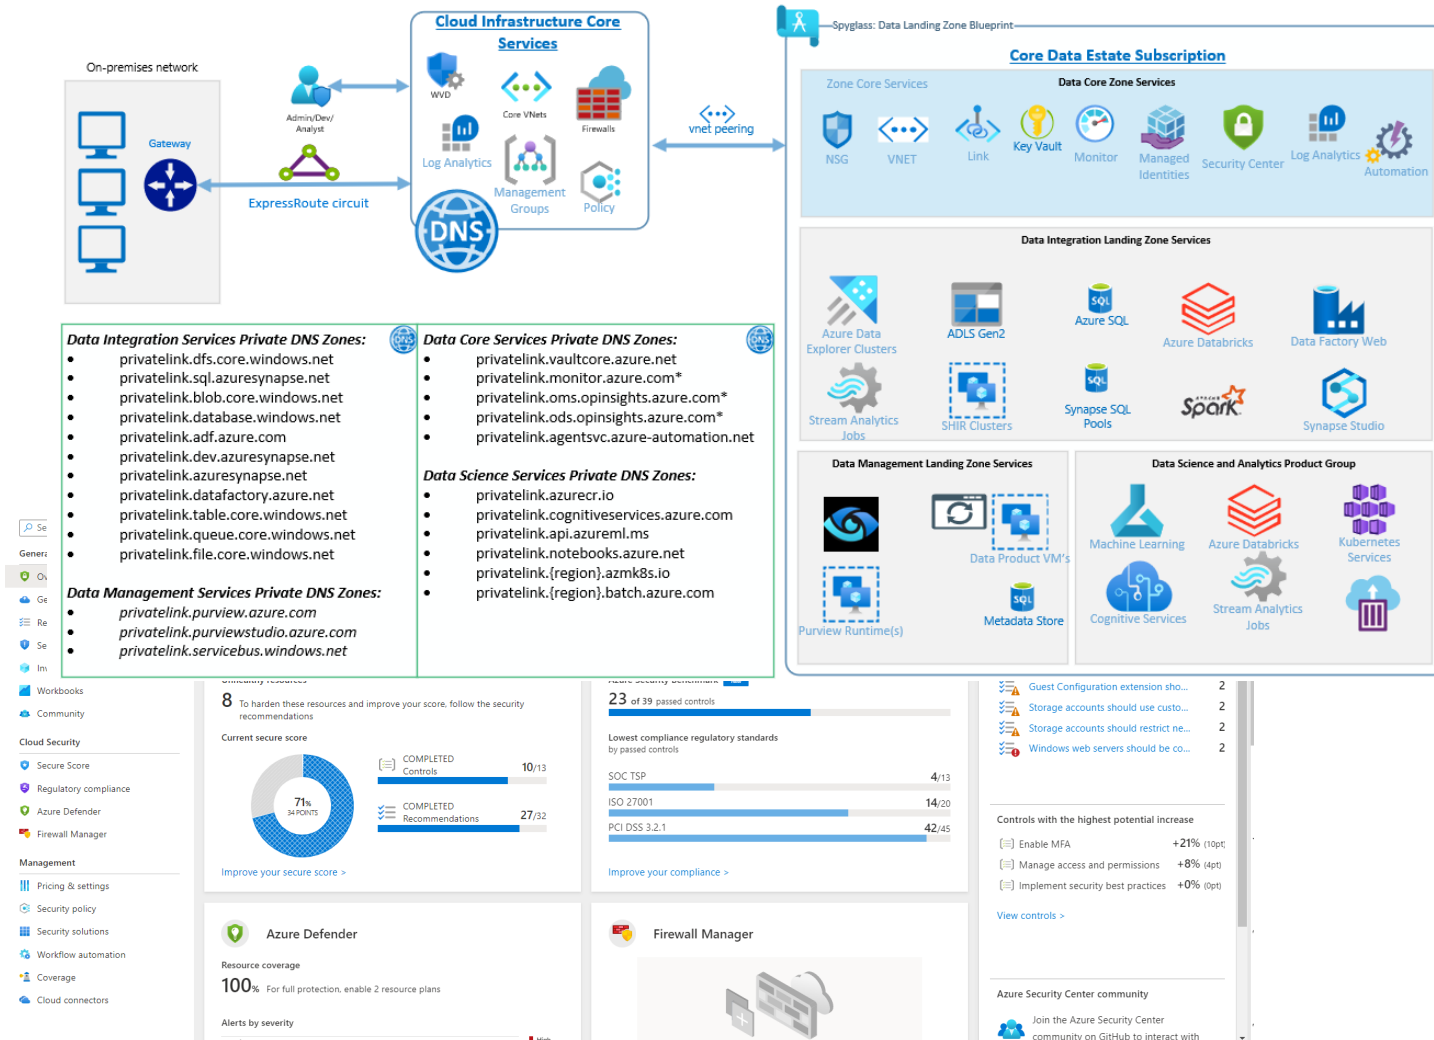
It is the foundation for migrating and modernizing existing analytics solutions, whether this be data dataliances or traditional data warehouses.

By establishing connectivity and integration, organizations can transform their unstructured and siloed data into a valuable strategic asset through:

- Data modernization backed by the Microsoft Azure Cloud
- Cloud native applications at any scale
- Responsible, powerful AI to make more informed decision-making
- Analytics and insights at a faster rate
- Responsible machine learning and artificial intelligence
- Governance backed by Microsoft Purview

Microsoft Fabric

17

| Data Factory | Synapse Data Engineering | Synapse Data Science | Synapse Data Warehousing | Synapse Real Time Analytics | Power BI | Data Activator |

OneLake

Unified
Product experience | Compute and storage | Governance and security | Business model

SPYGLASS MTG

Fabric is the SaaS Data Landing Zone

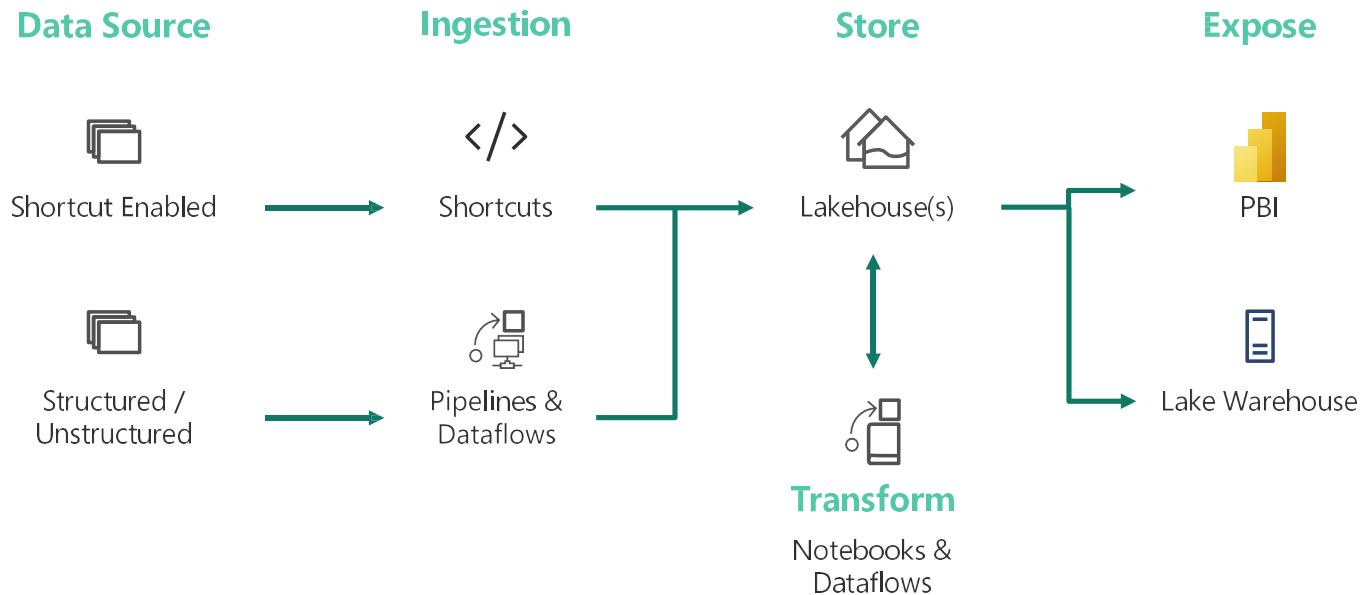| Layer	SaaS Fabric Capacity
Core services	• Network • Monitoring • Shared integration runtimes • CI/CD Agents • Security Agents
Core Data services	• Metastore for Data Services • Data Lake Services • Data Storage • Integration runtimes • Apache Spark • Synapse Analytics
Data Applications	• Machine Learning • Open AI • Cognitive Services • Data Share • Stream Analytics
Visualization	• Power BI • Data Flows

Lakehouse

The Microsoft Fabric Lakehouse analytics scenario makes it so that data can be ingested into OneLake with shortcuts to other clouds repositories, pipelines, and dataflows in order to allow end-users to leverage other data.

Once that data has been pulled into Microsoft Fabric, users can leverage notebooks to transform that data in OneLake and then store them in lakehouses with medallion structure.

From there, users can begin to analyze and visualize that data with Power BI using the see-through mode or SQL endpoints.

The Data Lakehouse scenario

Data Source	Ingestion	Store	Expose

Shortcut Enabled

Shortcuts

Lakehouse(s)

PBI

Structured / Unstructured

Pipelines & Dataflows

Lake Warehouse

Transform

Notebooks & Dataflows

Build and implement an end-to-end lakehouse for your organization:

1. Create a Microsoft Fabric workspace

2. Quickly create a lakehouse – an optional module to implement medallion architecture (Bronze, Silver, and Gold)

3. Ingest, transform and load data into the lakehouse – bronze, silver and gold zones as delta lake tables for medallion architecture

4. Explore OneLake, OneCopy of your data across lake mode and warehouse mode

5. Connect to your lakehouse using TDS/SQL endpoint

6. Create Power BI reports using DirectLake – to analyze sales data across different dimensions

7. Orchestrate and schedule data ingestion and transformation flow with Pipeline

8. Cleanup resources by deleting the workspace and other items
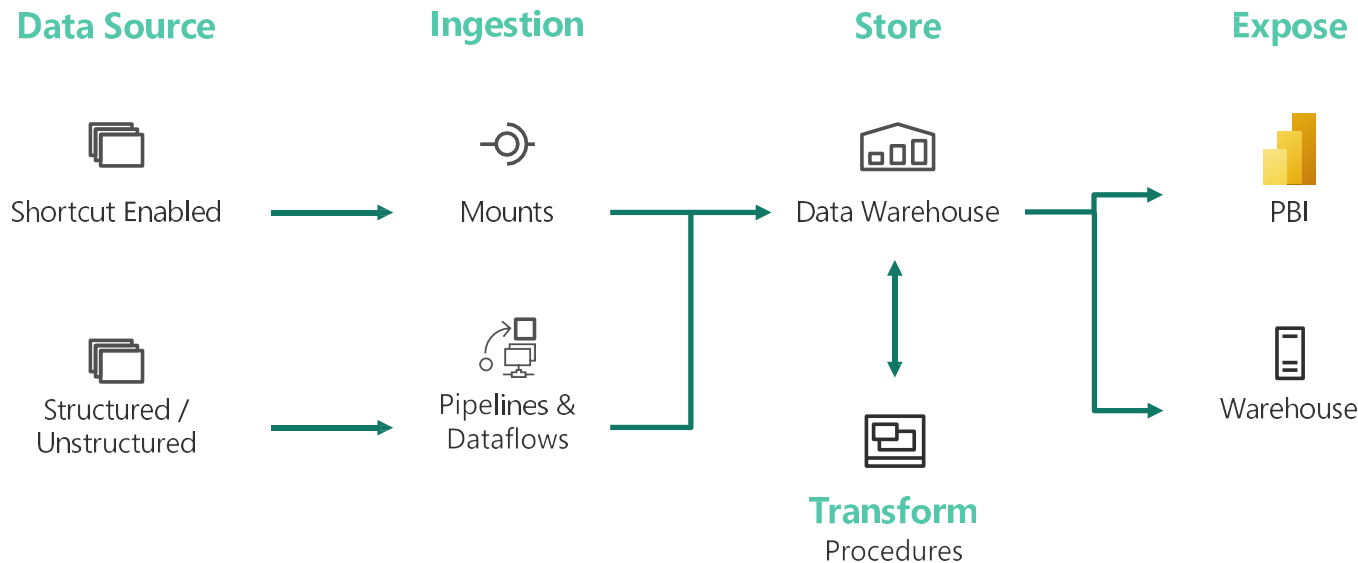
Experiment & Learn

SPYGLASSMTG

Data Warehouse

The Data Warehouse analytics scenario takes existing sources that are mounted, while pipelines and dataflows can bring in all other data that is needed.

IT teams can then define and store procedures to transform the data, which is stored as Parquet/Delta Lake files in OneLake.

From there, business users can analyze and visualize data with Power BI, again using the see-through mode or SQL endpoints.

The Data Warehouse scenario

Data Source — Shortcut Enabled / Structured / Unstructured

Ingestion — Mounts / Pipelines & Dataflows

Store — Data Warehouse

Transform — Procedures

Expose — PBI / Warehouse

Build and implement an end-to-end data warehouse for your organization:

1. Enable Microsoft Fabric in your tenant

2. Create a Fabric workspace

3. Quickly create a data warehouse

4. Ingest data from source to the data warehouse dimensional model

5. Transform the data to create aggregated datasets using T-SQL

6. Perform orchestration, data ingestion, and data transformation with pipelines

7. Query the data warehouse using T-SQL and a visual query editor

8. Create Power BI report using DirectLake mode to analyze the data in place

9. Cleanup resources by deleting the workspace and other items

Experiment & Learn

SPYGLASSMTG