

# Unlocking Transformative Data Value with Microsoft Fabric

A three-phased approach

# Forward

Today's world is awash with data—ever-streaming from the devices we use, the applications we build, and the interactions we have.

Organizations across every industry have harnessed this data to digitally transform and gain competitive advantages. And now, as we enter a new era defined by AI, this data is becoming even more important.

Generative AI and language model services, such as Azure OpenAI Service, are enabling customers to use and create everyday AI experiences that are reinventing how employees spend their time. Powering organization-specific AI experiences requires a constant supply of clean data from a well-managed and highly integrated analytics system. But most organizations' analytics systems are a labyrinth of specialized and disconnected services.

And it's no wonder given the massively fragmented data and AI technology market with hundreds of vendors and

thousands of services. Customers must stitch together a complex set of disconnected services from multiple vendors themselves and incur the costs and burdens of making these services function together.

That's why we've introduced a new innovation—Microsoft Fabric—an end-to-end, human-centered analytics product that brings together data and analytics tools in one place. Fabric integrates proven technologies like Azure Data Factory, Azure Synapse, and Microsoft Power BI into a single unified product, empowering data and business professionals alike to unlock the potential of data and lay the foundation for the era of AI.

## **The data platform for the era of AI: Introducing Microsoft Fabric**

# Contents

Forward	2
Current state of data, analytics, and business intelligence	4
The challenge: an organically evolved data estate	5
Complex data and analytics challenges	6
Now trending: data mesh, data fabric, data hub	7
Modern data architectures are not mutually exclusive. They are collectively transformative.	8
Introducing Microsoft Fabric – a unified analytics solution for the era of AI	9
A unified SaaS-based solution that stores all organizational data where analytics workloads operate	11
Three- phased approach to unlocking data value	12
<b>ENVISION</b>	<b>13</b>
The art of the possible	14
Healthcare	15
Financial Services	16
Public Sector	17
Retail	18
Sustainability/Energy	19

<b>EXPERIMENT &amp; LEARN</b>	<b>20</b>
Marketing	21
Operations	22
HR	23
Sales	24
Identify technical experts and business users for the POC	25
Document clear goals for your POC and the criteria that will define its success	26
Using common analytics scenarios and proposed activities to guide your Proof of Concept	27
Lakehouse	28
Data Warehouse	30
Data Science	32
Real time analytics	34

<b>BUILD &amp; SCALE</b>	<b>35</b>
Building a foundation of alignment	36
Aligning with your organization’s leaders and data stewards	37
Conversations and questions to drive your business value roadmap.	38
Conversations and considerations to drive your infrastructure roadmap	40
Mapping your current data and analytics estate	41
Ideate on planned transformation and outcomes	42
Establishing a network of transformation champions	43
Training your network of champions	44
You are on your way to transformative data value	45

# Current state of data, analytics, and business intelligence

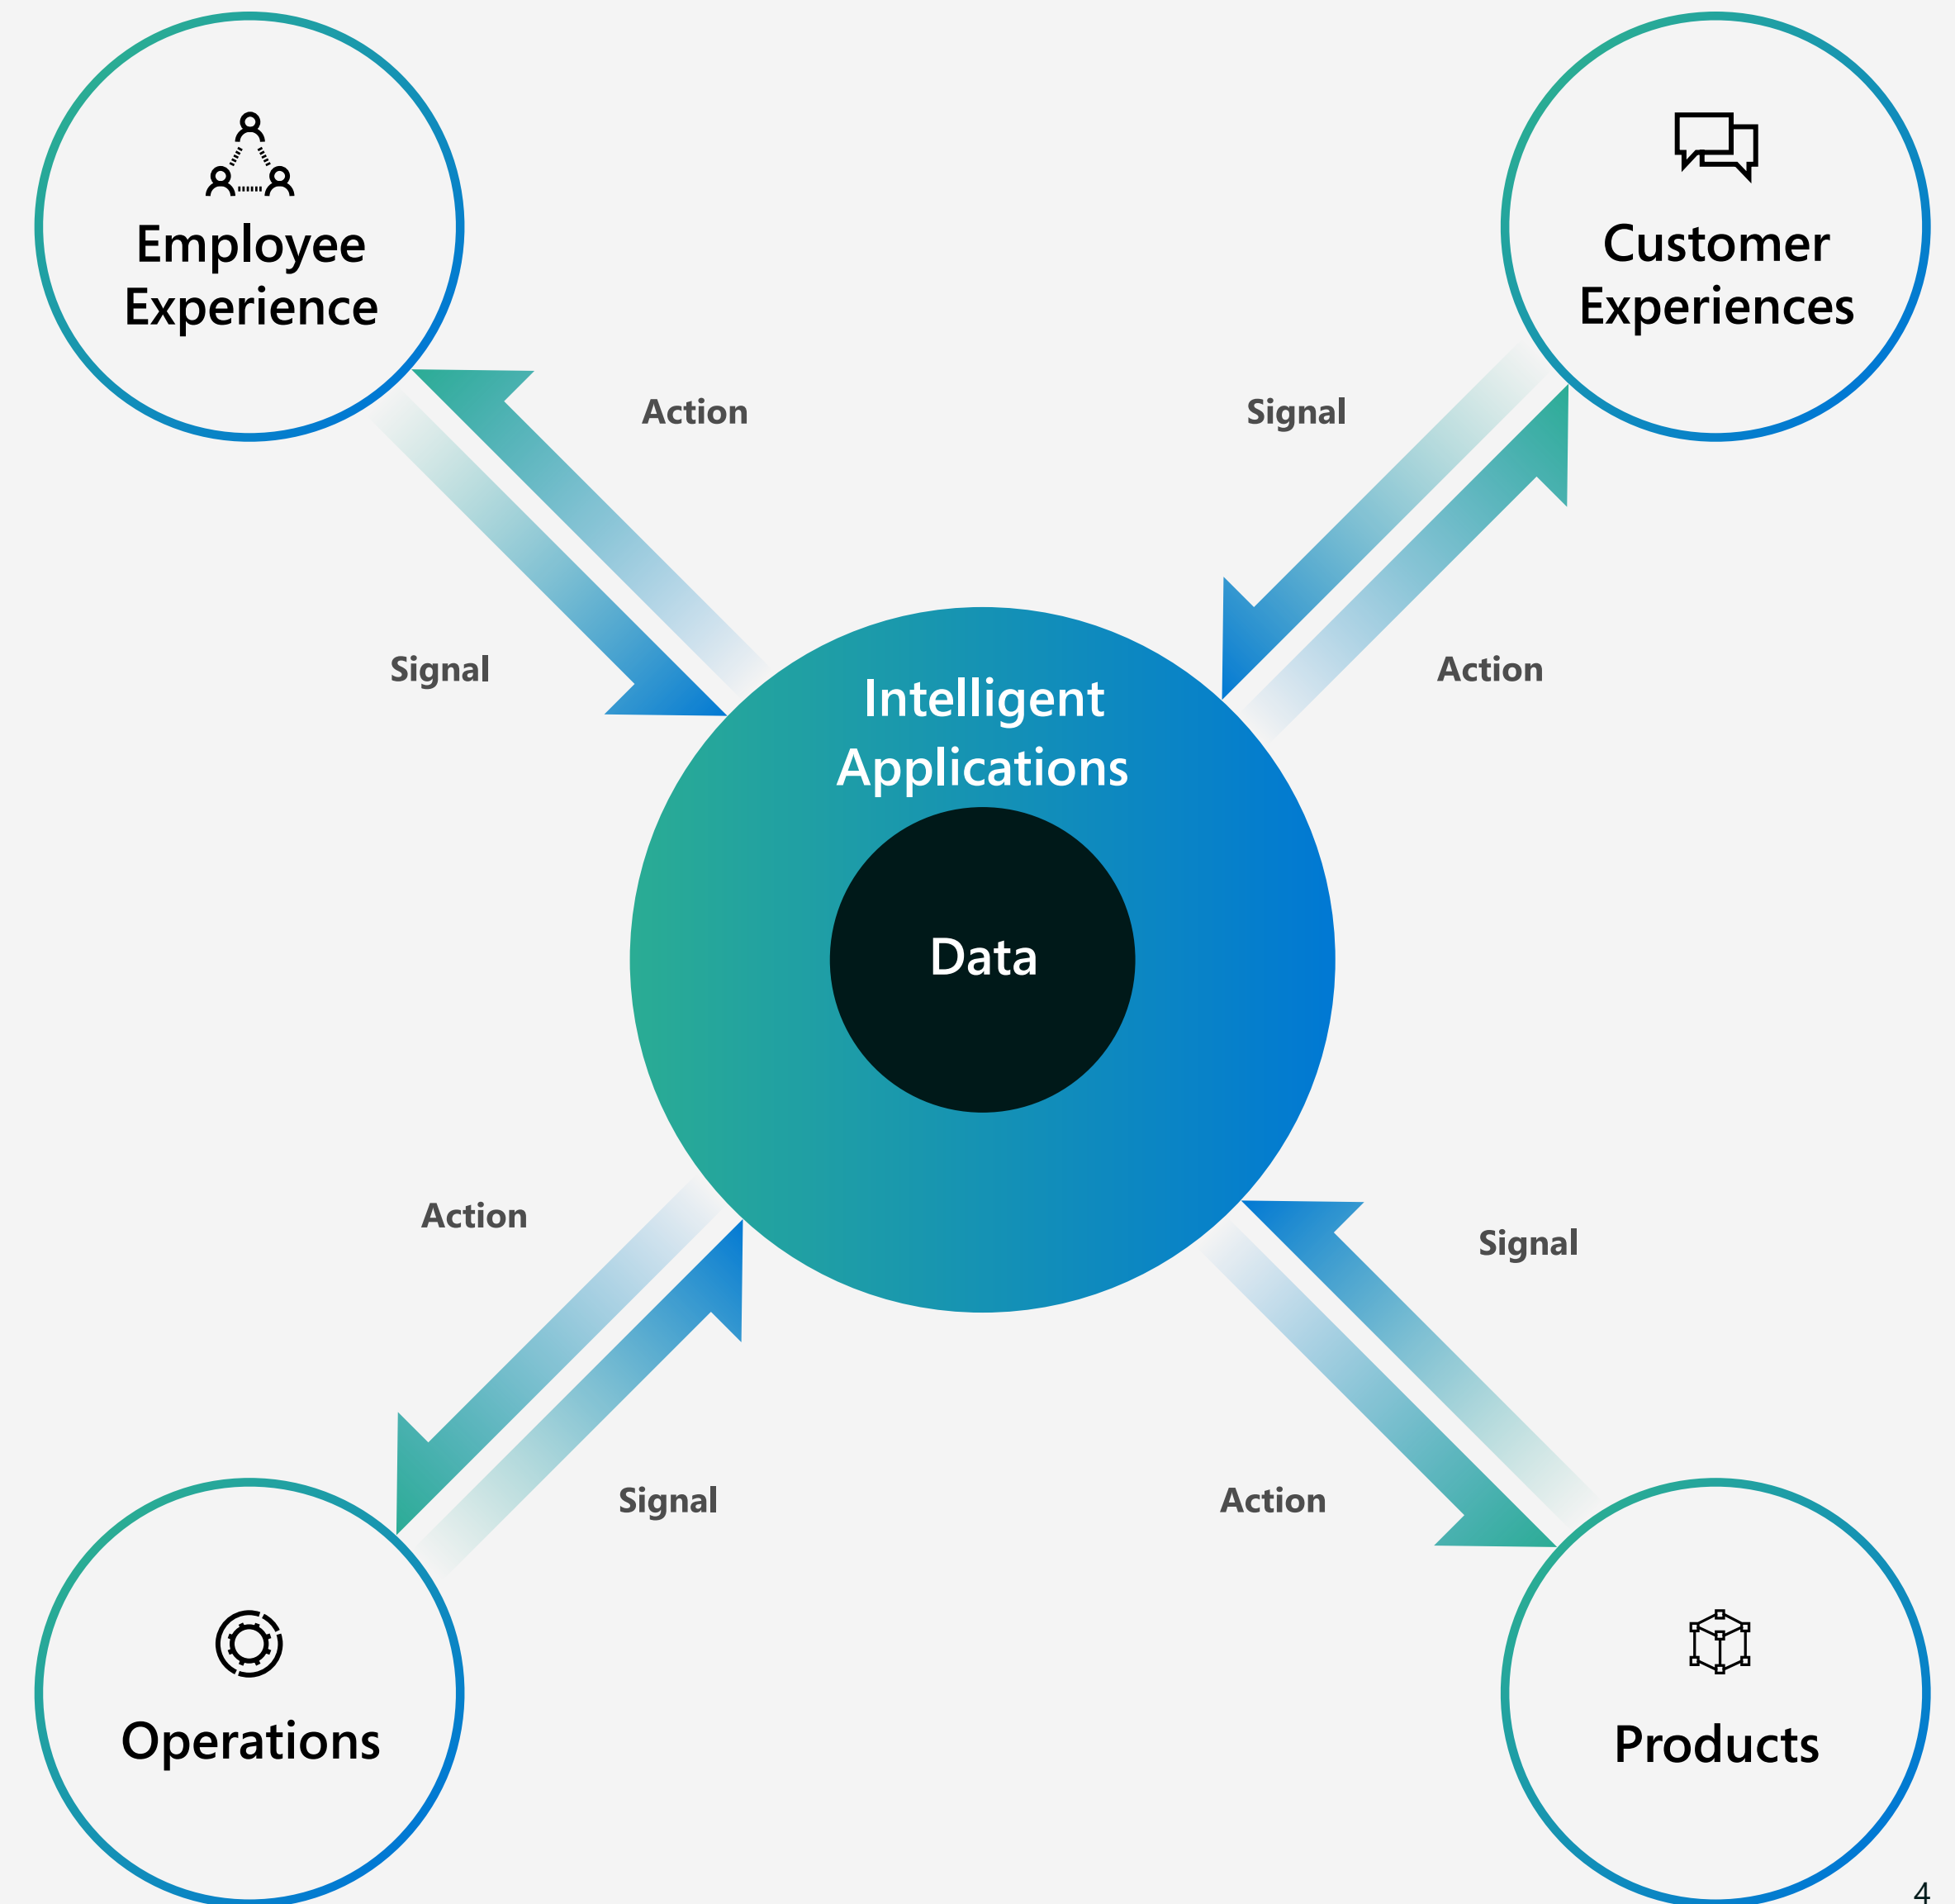
## Data has become the oxygen of digital transformation

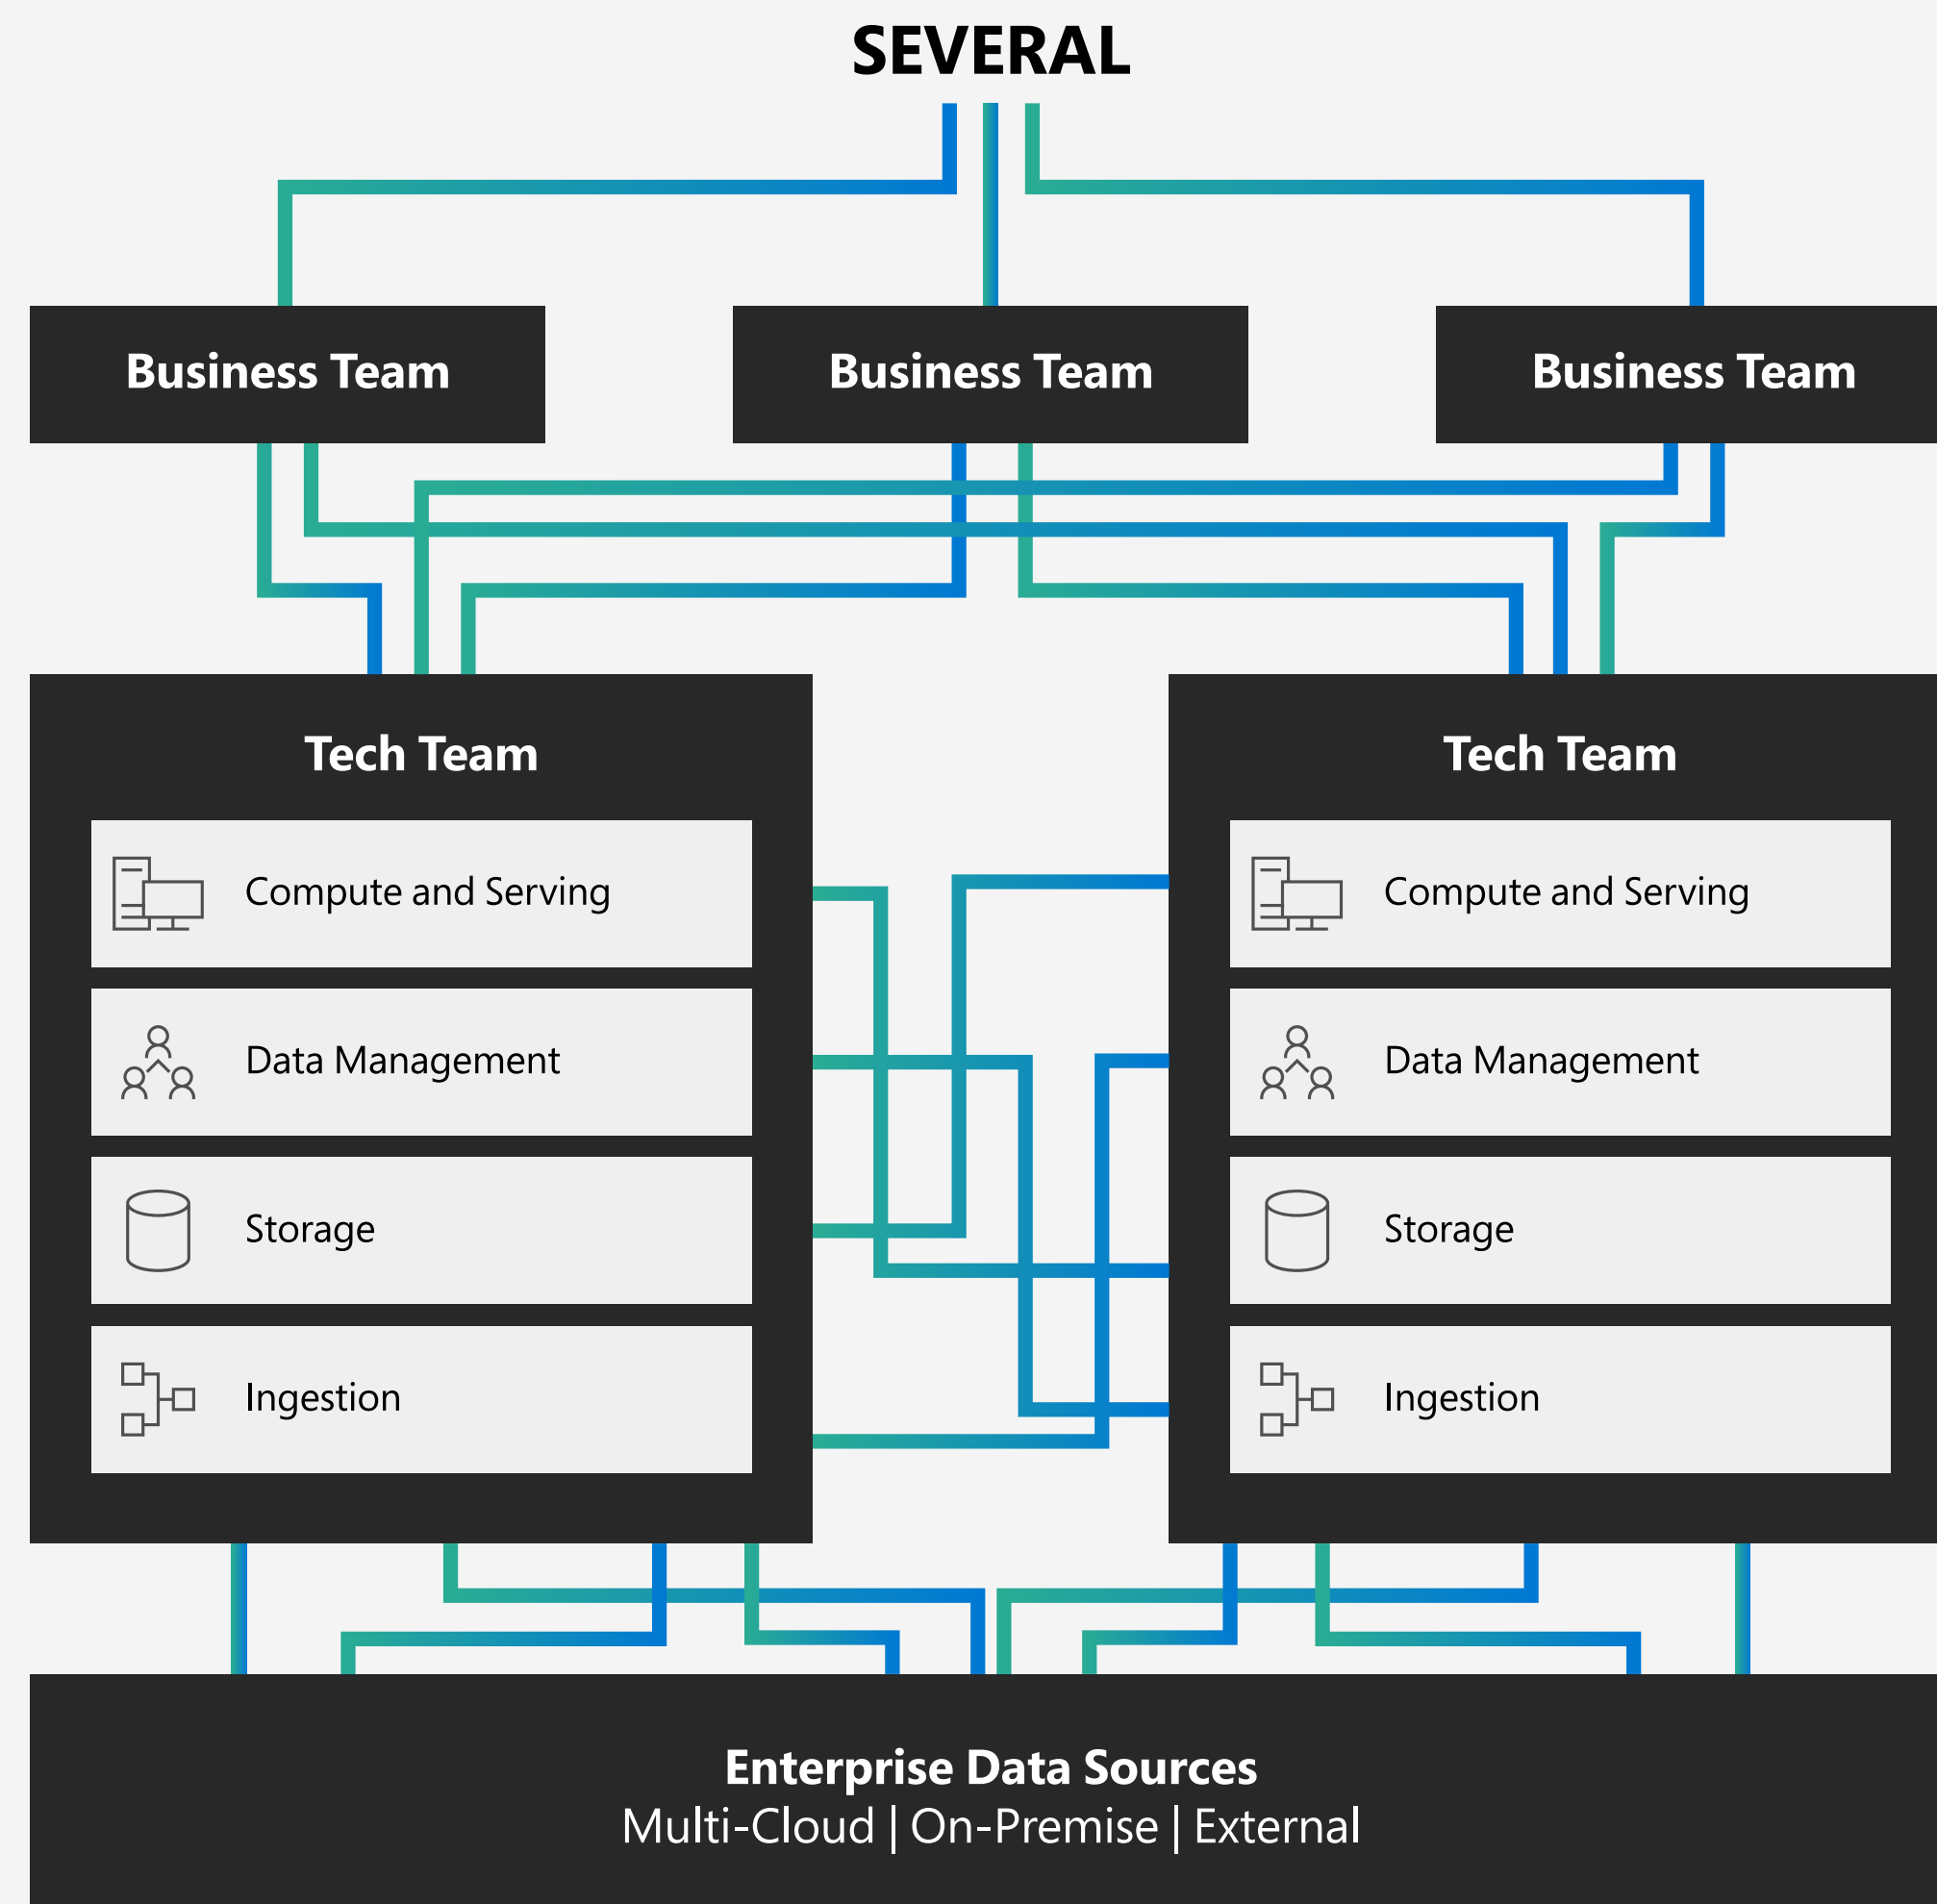
Data is now truly driving digital transformation: from creating delightful customer experiences and enhancing employee productivity to generating operational efficiencies. But it's not just the raw data itself. It's the integrated insights and intelligence that can be activated across an entire connected organization. Unlocking the value from your data and bringing it into the era of AI is the key to driving innovation within your organization and for your customers.

- Personalized Employee Experiences
- Smart Facilities
- Actionable Seller Intelligence
- Personalized Support Experiences
- Finance Intelligence
- Digital Security and Risk Intelligence

Yet, research shows that only **32% of organizations reported being able to realize tangible and measurable value from data.** And that is where the challenges begin to reveal themselves.

Source: [Accenture](#)





# The challenge: an organically evolved data estate

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

# Complex data and analytics challenges

Organizations know that to succeed and thrive in today's constantly changing landscape, they must be able to transform the massive amount of data they have into deep, impactful insights. We've talked to a wide range of customers who have led their organizations in data-first modernization. And we understand that businesses today face a variety of data and analytics challenges that can be difficult to navigate:

- **Scaling data and analytics** across the organization while reducing costs and optimizing existing data and management
- **Gaining BI adoption** to streamline data sharing and agility to insights across Lines Of Business and teams
- **Encouraging data literacy** by making data more accessible and easier to handle across technical and non-technical users


- **Difficulty in enabling data-driven organization** due to siloed data and lack of unified tools
- **Balancing the need for data access and self-service** analytics while remaining governed
- **Limited scalability** of legacy solutions as data demand rises exponentially
- **Breaking down data siloes** across the organization into a unified source of truth
- **Delivering on the promise of analytics** with limited resources

Investing in a comprehensive data and analytics platform can address the hurdles encountered in your current data ecosystem. With higher adoption of cloud-based business intelligence and analytics tools, it's imperative to work with a partner that has a trusted portfolio of security products—that can enable faster, data-driven decisions in a secure manner.



**55% of companies** have a mostly manual approach to discovering data within their enterprise.

Source: [Accenture](#)



**32% of companies** reported being able to realize tangible and measurable value from data.

Source: [Accenture](#)

# Now trending: 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.

# Modern data architectures are not mutually exclusive. They are collectively transformative.



We know that “by 2024, **75% of organizations will have established a centralized data and analytics (D&A) center of excellence** to support federated D&A initiatives and prevent enterprise failure.” Federated—otherwise known as standardized—data is a critical component of creating a centralized data and analytics (D&A) center of excellence, which is increasingly becoming a key priority for organizations. By implementing a standardized approach to managing data, organizations can ensure that their data is consistent, accurate, and reliable, making it easier to analyze and derive insights from.

Source: [Gartner, Our Top Data & Analytics Predicts](#)

While there is copious academic debate surrounding these modern architectures, the reality is the combination of them will drive the best solution.

**The Microsoft approach** is based on best practices refined over more than two decades of firsthand experience in building global-scale products and services. It starts with a holistic view of the organization, factoring in the people, processes, culture, and technology. Then, by applying data governance, security, and compliance across every layer of the stack, it ensures customers have a truly innovative environment that empowers everyone to do their best work.




# Introducing 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.



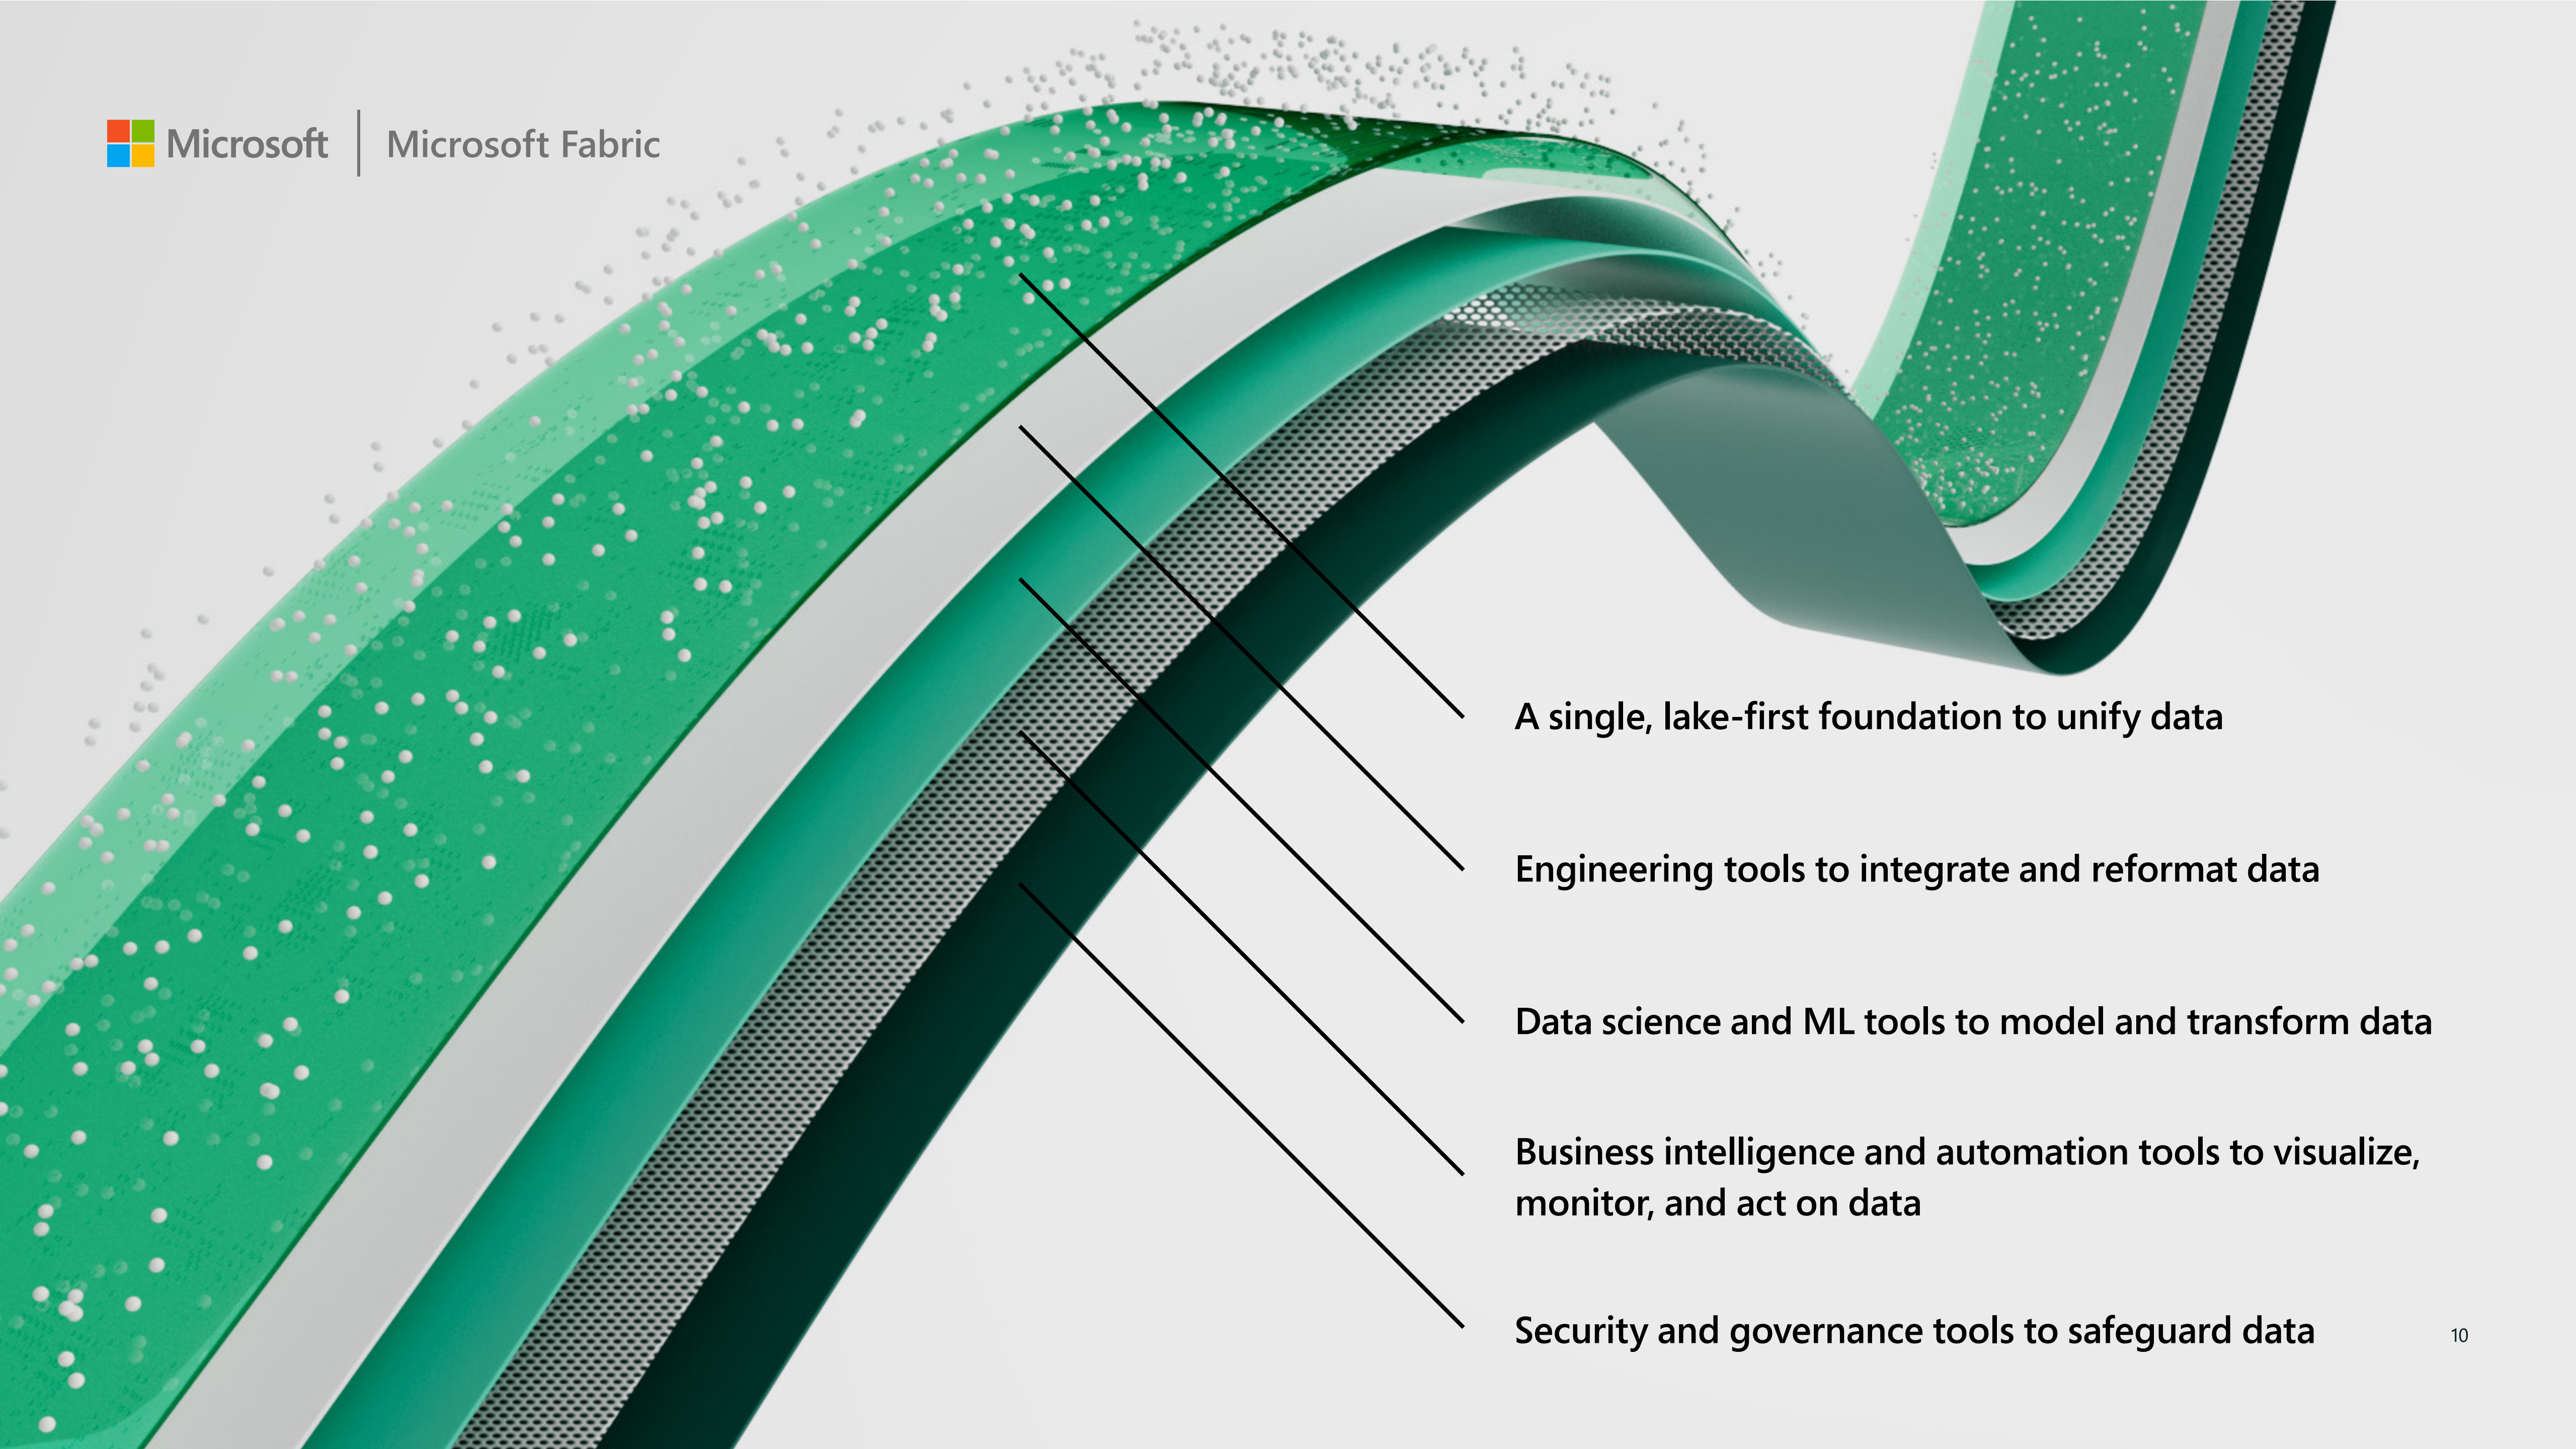
## Business Intelligence

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.

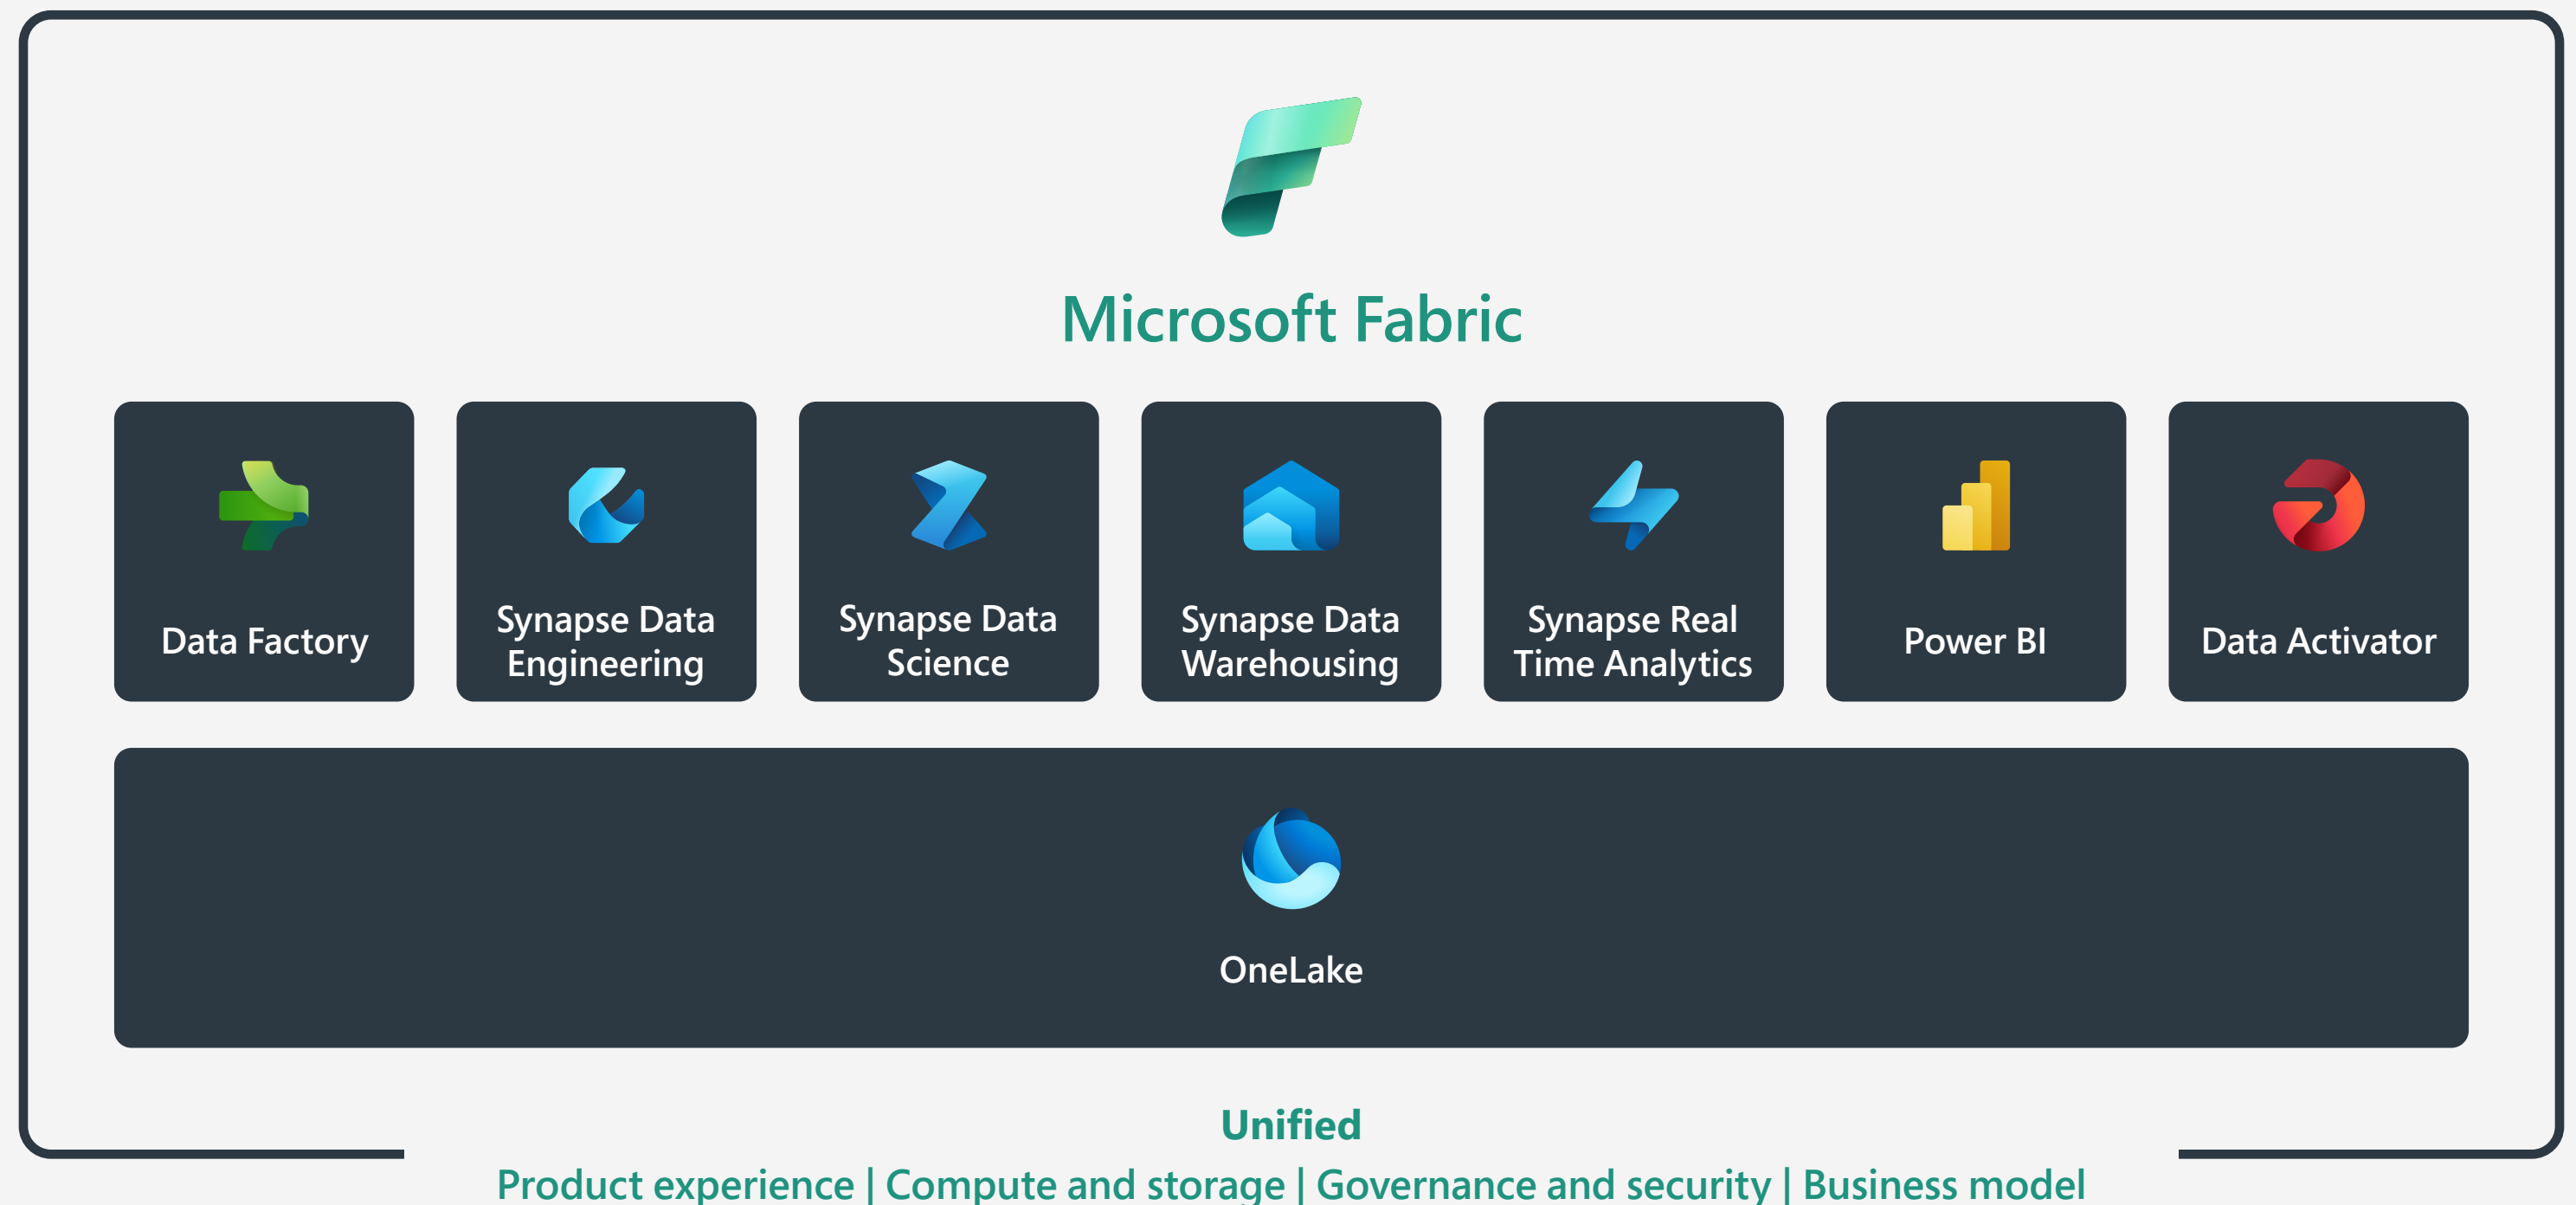


# A unified SaaS-based solution that stores all organizational data where analytics workloads operate

**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 appliances 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



# Three-phased approach to unlocking data value

## ENVISION

Learn about the opportunities that are possible for your organization when you transform your data estate to unlock the value of your data. And how Microsoft Fabric makes it all possible.

## EXPERIMENT & LEARN

Begin to develop data culture and build trust through a smaller scale Proof of Concept that delivers immediate business value.

## BUILD & SCALE

Get aligned with your organization's leaders and data stewards on the business value you'll continue to drive over time, map your current and future data estate, and establish a network of transformation champions.

What will you empower your organization to achieve? What innovation will come from transformation of your data estate and acceleration of the value you derive from your data?


This guide breaks down your journey into three phases: Envision, Experiment & Learn, and Build & Scale. As we walk through each of the phases, we will help you discover and consider all the factors involved in aligning teams with a unified source of truth, establishing a thriving center of enablement for every user with secure, democratized insights across your organization, and unlocking transformative data value.


PHASE 1

# Envisioning the opportunity begins with customer centricity

To begin envisioning what innovation may come from unlocking your data's value, it's important to start with understanding the motivations, challenges and experiences of your employees, customers, partners, and suppliers.

By grounding your data estate transformation initiative in the outcomes that are made possible for all constituents – you will help keep customer centricity at the forefront during the Build and Scale phase.





## PHASE 1: ENVISION

# The art of the possible

What are your customers, employees, partners and suppliers' unmet needs and how can you uniquely solve those unmet needs through unified data, comprehensive insights and faster decision making? Let's have a closer look at what Microsoft Fabric makes possible for the following industries.



# Healthcare

The healthcare industry comes with many challenges, as it handles massive amounts of highly sensitive and regulated data. In addition to maintaining high levels of compliance and security, Microsoft Fabric allows for an improved patient experience by giving healthcare organizations access to structured and unstructured medical data in OneLake. Having a lakehouse foundation enables healthcare professionals to run analytics to identify patient needs in real-time.

Healthcare organizations can also leverage real-world data collected by clinical drug trials. By using machine

learning and artificial intelligence, coupled with large sets of genomic and population data, healthcare professionals can gain new insights into drug efficacy. Healthcare providers will also be empowered to collaborate more closely by being able to share and access patient data and insights.

With Microsoft Fabric, healthcare can be reimaged by analyzing massive amounts of data at a faster pace than ever before. Paired with improved quality of insights, healthcare professionals are equipped to curate personalized treatment plans for all patients.

## Data challenges

- No comprehensive source for all aggregated data
- Securely providing care teams access to the data they need
- Slow time to insights

## How Microsoft Fabric can help

- Leverage OneLake to aggregate all data into a single source of truth
- Empower health team collaboration through democratized insights
- Analyze massive amounts of data at an accelerated pace, using built-in security and governance policies

## Key outcomes

- Enable decision-making backed by data to reshape care and insights
- Improve the patient experience by creating an accessible, holistic view across healthcare professionals
- Deliver secure, real-time clinical and operational insights using Event Hub, IoT Hub, Kafta, and more



# Financial Services

The financial services industry (FSI) is heavily regulated and governed, which makes being backed by the security of a Microsoft product that much more important. With Microsoft Fabric, FSI professionals can access customer portfolio and market data with open and governed access controls. They can also strengthen risk prevention by leveraging scalable compute and analytics power modeling.

In the banking industry, organizations will be able to drive increased customer lifetime value and loyalty by providing a holistic view of the customer's current

financial state, allowing them to upsell or cross-sell products and services that best fit those individual customer needs.

In the insurance industry, organizations will be able to gain insights about customers by connecting internal and external data to provide custom solutions to the customer.

In capital markets, organizations can initiate more personalized client interactions by gaining more insights into a client's profitability and predictive analytics, enabling them to sell the right product at the right time to close more deals.

## Data challenges

- Abundance of data and technology siloes
- Securely collecting, storing, and sharing data in a highly regulated industry
- Slow time to informed decision-making and risk responses

## How Microsoft Fabric can help

- Collect data in OneLake to provide one unified source of truth
- Access all portfolio, reference, market, and risk data with open and governed access controls
- Analyze all different types of data in real-time, using built-in security and governance policies

## Key outcomes

- Enable decision-making backed by data to thrive in a competitive environment
- Gain a complete, 360-degree view of customers through secure data sharing
- Drive accurate insights with speed, while strengthening risk detection and prevention





# Public Sector

For most government agencies, data is scattered. Data cannot be accessed digitally, and it is not easily digestible. With Microsoft Fabric, customers will be able to collect and store data in one single location—and on the cloud.

Government agencies will now be able to combine research and data across the public health sector and use machine learning and artificial intelligence to identify health risks and trends.

In addition, governments can leverage data to drive improved urban planning by offering better services and experiences. Streaming data from IoT

sensors and public utilities/transportation machinery can be ingested in OneLake with IoT Hub, for real-time analysis.

From a tax fraud prevention standpoint, governments will be able to develop advanced machine learning models to analyze large amounts of tax data being held in OneLake to identify and prevent tax fraud.

## Data challenges

- Abundance of data and technology siloes
- Securely collecting, storing, and sharing data in a highly regulated industry
- Slow time to informed decision-making and risk responses

## How Microsoft Fabric can help

- Collect data in OneLake to provide one unified source of truth
- Access all portfolio, reference, market, and risk data with open and governed access controls
- Analyze all different types of data in real-time, using built-in security and governance policies

## Key outcomes

- Enable decision-making backed by data to thrive in a competitive environment
- Gain a complete, 360-degree view of customers through secure data sharing
- Drive accurate insights with speed, while strengthening risk detection and prevention



# Retail

Retailers face their own set of unique challenges when it comes to data. They often struggle with closing the gap between online and physical stores, integrating data in a meaningful way, and efficiently managing supply chains. With Microsoft Fabric, retailers can leverage real-time data to inform their merchandising strategy, leveraging customer data to learn what target audiences value and curating personalized experiences.

Retailers can also optimize digital spend by creating models that identify who and where to target buyers based on data trends. Tracking store activity by

ingesting real-time data from in-store sensors and cameras will identify even more in-person areas for optimization.

Lastly, and most likely top-of-mind for many retailers, is being empowered to deliver an intelligent supply chain. By using data to anticipate customer demand on a local and global scale, retailers can get the right product, at the right price, to the right customer, at the right time—closing the gap between demand and fulfillment.

## Data challenges

- Siloed, incomplete data with complex architecture
- Using data to personalize customer experiences
- Meeting data regulations in areas where your organization does business

## How Microsoft Fabric can help

- Unify data from numerous sources such as purchase orders, inventory, and manufacturing all in a OneLake
- Gain a 360-degree view of all customers by collecting and analyzing site, digital, and smart store behavior
- Leverage built-in security and governance to help protect collected customer data

## Key outcomes

- Inform merchandising and supply chain strategy with real-time analysis and analytics
- Develop content, copy, and products tailored to your customers' specific tastes and interests
- Be rest assured that in-store and Ecommerce transaction data is secure



# Sustainability/Energy

With Microsoft Fabric, renewable energy providers can combine real-time data from wind turbines/solar panels, as well as customers, to identify future power demands through artificial intelligence and machine learning models. Energy providers will be able to collect large amounts of data from in-home systems and store it in OneLake for analysis.

Microsoft Fabric will enable energy providers to get more out of their investments through IoT-enabled remote tracking to drive AI analytics for sustainable decision-making. Energy providers will also be able to collect real-time data in OneLake to improve efficacy

of energy sources based on usage trends.

Most importantly, energy providers can reimagine energy as we know it by identifying new business models through AI-driven analysis and real-time consumption/demand.

## Data challenges

- Using data to transition to clean, renewable energy
- Forecasting and predicting energy demand
- Identifying new business models backed by data

## How Microsoft Fabric can help

- Collect real-time data in OneLake to help identify how clean energy is used and how to improve its efficacy
- Combine real-time data from disparate sources such as wind turbines and solar panels and use ML and AI to identify demand
- Identify new business models through AI-driven analysis of energy distribution, consumption, and customer demand

## Key outcomes

- Manage smart and efficiency energy on local and global scales
- Identify how much demand for power will be needed in the future – quickly and efficiently
- Drive efficient, sustainable best practices and reduce environment impact




## PHASE 2

# Experiment & Learn

The best first step is to experiment and learn through a Proof of Concept (POC). Starting with a departmental scenario rather than an organization-wide scenario can help you ensure that time spent is minimized and the value derived is maximized.

Consider the following departmental scenarios as you embark on your POC.



# Marketing




## Improve campaign analysis and planning

- Unify campaign data from different sources into a single location to make it easier to track campaign performance vs targets and budget.
- Purchase and attribution data can be brought in to develop models that help identify what demographics and/or areas to target for future campaigns.
- Translate data into action by receiving real-time alerts when important campaign metrics change and adjusting if needed.
- Pull data silos together by compiling information in a single, unified view across your marketing technologies—analyze multiple campaigns together, compare data across markets and platforms, and more.
- Track performance vs. targets and budget, using fast, interactive reports to help show ROI and manage spend.



## Optimize paid media spend

- Impressions and sales data can be combined with media spend and attribution data in OneLake.
- Develop models that help identify the best areas and channels to move budget to help gain the most impressions or leads per dollar.



## Streamline location analytics

- Track retail brand awareness campaign effectiveness at the regional and individual store level by collecting store and behavioral data in one centralized lake storage.
- Analyze this data with advanced methods and present complex behavioral data in visually stunning, digestible formats.
- Improve understanding of customer purchase patterns by combining digital campaign data with in-store visit and Ecommerce data into a single view.
- Target customers more effectively, pulling from cloud and on-premises data sources to see footfall patterns and time specials according to when your target demographic is shopping.



## Improve website, social, and email analytics

- Collect real-time data like impressions, click-through rates, and video completion rates from digital ads on websites, social channels, and emails in one single location.
- Perform analysis by analyzing product, region, and segment level performance to develop better targeted ads for customers you want.

# Operations



## Undergo production optimization

- Collect and analyze data from different manufacturing centers to identify production bottlenecks before they make an impact on the business.
- Develop predictive models to help them understand when to scale production up and down based on market trends. Increase efficiency by identifying production bottlenecks faster using a heads-up dashboard and custom diagrams.
- Use visualizations to easily understand data from manufacturers or production lines in one dashboard to aid in planning.
- Understand anomalies in real-time manufacturing data using drill-down capabilities, which better informs required corrective action.



## Enhance inventory planning and performance

- Unify sales data along with production data to identify which products to prioritize based on trends in the market and address stock levels before inventory shortages hit.
- Analyze which products to prioritize by easily spotting trends, such as regional spikes in demand or increased inventory cycle times for certain items.
- Assess and adjust levels of safety stock by combining recent trend and historical data from cloud, on-premises, or hybrid data sources.
- Receive notifications of major changes in available inventory across different geographies by configuring mobile alerts to help avoid inventory shortages.



## Improve sales and operations planning

- Merge sales, marketing, and manufacturing data from different clouds and systems of record to draft new product development or production plans more easily.
- Improve sales and production forecasting quality by analyzing metrics like past total sales, production output and shipment times together using natural language querying or other advanced techniques to move from hindsight to insight.
- Get better visibility into demand and supply projections across departments from anywhere with all data stored in the same logical data lake.



## Streamline fulfillment and distribution

- Make it easier to identify over- and under-used resources using production pick rates to predict whether to flex up or down and assess distribution center efficiency.
- Notify your customers sooner when there is a shipping delay by leveraging alerts from dashboards that can inform you about a shipping gap.

# HR



## Reduce employee turnover and gain visibility into recruiting and retention

- Combine data from financial apps and recruiting systems into custom reports in just minutes, leveraging rich visualizations that drive faster time to insight.
- Explore your staffing funnel or recruiter performance with an intuitive search-engine format that lets you ask questions of your data and get instant, visual answers.
- Stay informed of fluctuations in key metrics, like when retention rates fall below a certain threshold or key positions are filled by setting up mobile notifications.



## Facilitate workforce/demand planning

- Combine all your workforce data from Dynamics 365 Human Resources into a single dashboard for easier viewing and manipulation to drive greater insight into your business.
- Get important updates quickly, such as unplanned changes to budget or headcount, with real-time data refresh and custom mobile alerts.
- Investigate demand planning anomalies quickly and easily by clicking into detailed reports directly from your dashboard, or by leveraging the search-engine-like capability to ask questions of your data and get immediate answers.



## Monitor benefits, rewards, and compensation

- Combine disparate HR data in one location by seamlessly connecting to both on-premises and cloud-based data sources, such as Human Capital Management systems like Dynamics 365 Human Resources and recruiting platforms like LinkedIn Recruiter.
- Report on the rewards and compensation metrics you care about by creating visualizations and custom dashboards instantly, instead of being locked into pre-built vendor solutions.
- Investigate unexpected spikes in benefit usage or other areas of interest by instantly diving into underlying reports directly from your dashboard and develop models that can help predict when these spikes may be expected in the future.



## Predict organizational health trends

- Predict health trends, such as the type and frequency of sickness, enabling companies to prepare for outbreaks.

# Sales



## Identify better opportunities for upselling and cross-selling

- Bring together various data points into a bigger, more detailed picture to reveal repeat sales patterns, influences on repeat sales journeys, and where customers are in their journey.
- Identify upsell and cross-sell opportunities and recommended actions to capture those opportunities.



## Increase localization

- Tailor sales efforts to local markets to uplift and maximize overall sales.
- Use sales analytics to understand sales conversion at a micro level to structure sales initiatives, promotional campaigns, and even stocking levels that best fit local markets.




## Develop enhanced pricing management plans

- Build a thorough price-management plan to use detailed analytics on sales, promotions, and competitors.
- Aggregate this data into a single pane of glass to help teams quickly slice and analyze data in various forms to deliver better pricing and promotion strategies.



## Improve sales performance analysis

- Monitor and compare key metrics from different data sources in real time, including win rates, number of wins, revenue to plan, margin, discounts, and more.
- Develop a single dashboard that makes it easy for anyone to access and track these metrics on their own.
- View actuals vs. forecast by territory, product line, and customer segment, and get answers to questions like “Which territory had the highest win rate last month?” in the form of interactive charts.





# Sales




## Drive more visibility into pipeline management:

- See your pipeline and quota in one view by connecting data from to your CRM account, forecasting tools, and more in one data lake.
- Get instant insight into key metrics like pipeline by sales stage, average deal size, and pipeline velocity so you can focus time and energy on the right opportunities.
- Combine current and historical data from multiple sources, including Dynamics 365 Sales, to identify trends and more effectively qualify opportunities in or out.



## Improve quota management:

- Blend sales data from different sources and deliver real-time insights in dashboards and reports that make the latest quota attainment results available on demand. Keep close tabs on metrics like pipeline value and win rates to get a holistic picture of progress towards quota, and quickly identify at-risk territories.
- Manage team performance with personalized dashboards that enable drill-down into variance to target by individual representative.
- Define and develop models that help you plan for more accurate and attainable quotas in future quarters.



## Uncover trends:

- Increase deal size via market basket analysis—compile data from multiple sources to discover products frequently purchased together and identify cross-sell and up-sell opportunities.
- Analyze sales more rapidly to improve forecasting quality, identify spending trends and variations (e.g., by industry or region), and find new opportunities.
- Combine transactional information from multiple cloud and on-premises sources to uncover trends in top accounts by account manager, territory, product line and other factors.




## Discount analysis:

- Compare estimated revenue and commission against targets for specific categories, geographies, and account types.
- Measure the impact of discounts and promotions on demand, traffic, average sale amount, and other sales KPIs.
- Leverage historic data collected in legacy systems to analyze past discount trends and make better decisions on what to discount today.
- Improve decision making by modeling dynamic discount scenarios using what-if analysis and Quick measures.



# Identify technical experts and business users for the POC

Identify availability of technical experts and business users to support you during the POC execution. Ideally, the technical decision makers in your organization, either an IT leader and/or a data leader, will start by creating clear requirements for the proof of concept. Next, your data practitioners such as architects, engineers, data scientists and analysts will execute the POC. And lastly business decision makers or lines of business (LOB) leaders will play a large role in evaluating the results.



## WORKSHEET

# Document clear goals for your POC and the criteria that will define its success

Keep in mind that a POC should be a short and focused effort to quickly prove a limited set of concepts and capabilities. These concepts and capabilities should be representative of the overall workload. Start by creating goals for your POC.

Use your goals to identify specific tests and to provide the outputs you identified. It's important to make sure that you have at least one test to support each goal and expected output.

Goal	Output	Test 1	Test 2

### A few suggested goals

- Reduced number of products used from X to Y
- Time saved from X to Y
- Improve development time from X to Y
- Expand analytical potential between X and Y through the addition of the new data
- Expand user base from X to Y
- Satisfy user base demonstrated by X score on user surveys
- Improve dashboard performance from X to Y
- Decrease time to insight from X to Y
- Business value added. This should be specific to the scenario you choose
- Performance benchmarking
- Low code/no code data pipelining

# Using common analytics scenarios and proposed activities to guide your Proof of Concept

Microsoft Fabric supports four common end-to-end analytics scenarios:

1. Lakehouse
2. Data Warehouse
3. Data Science
4. Real Time Analytics

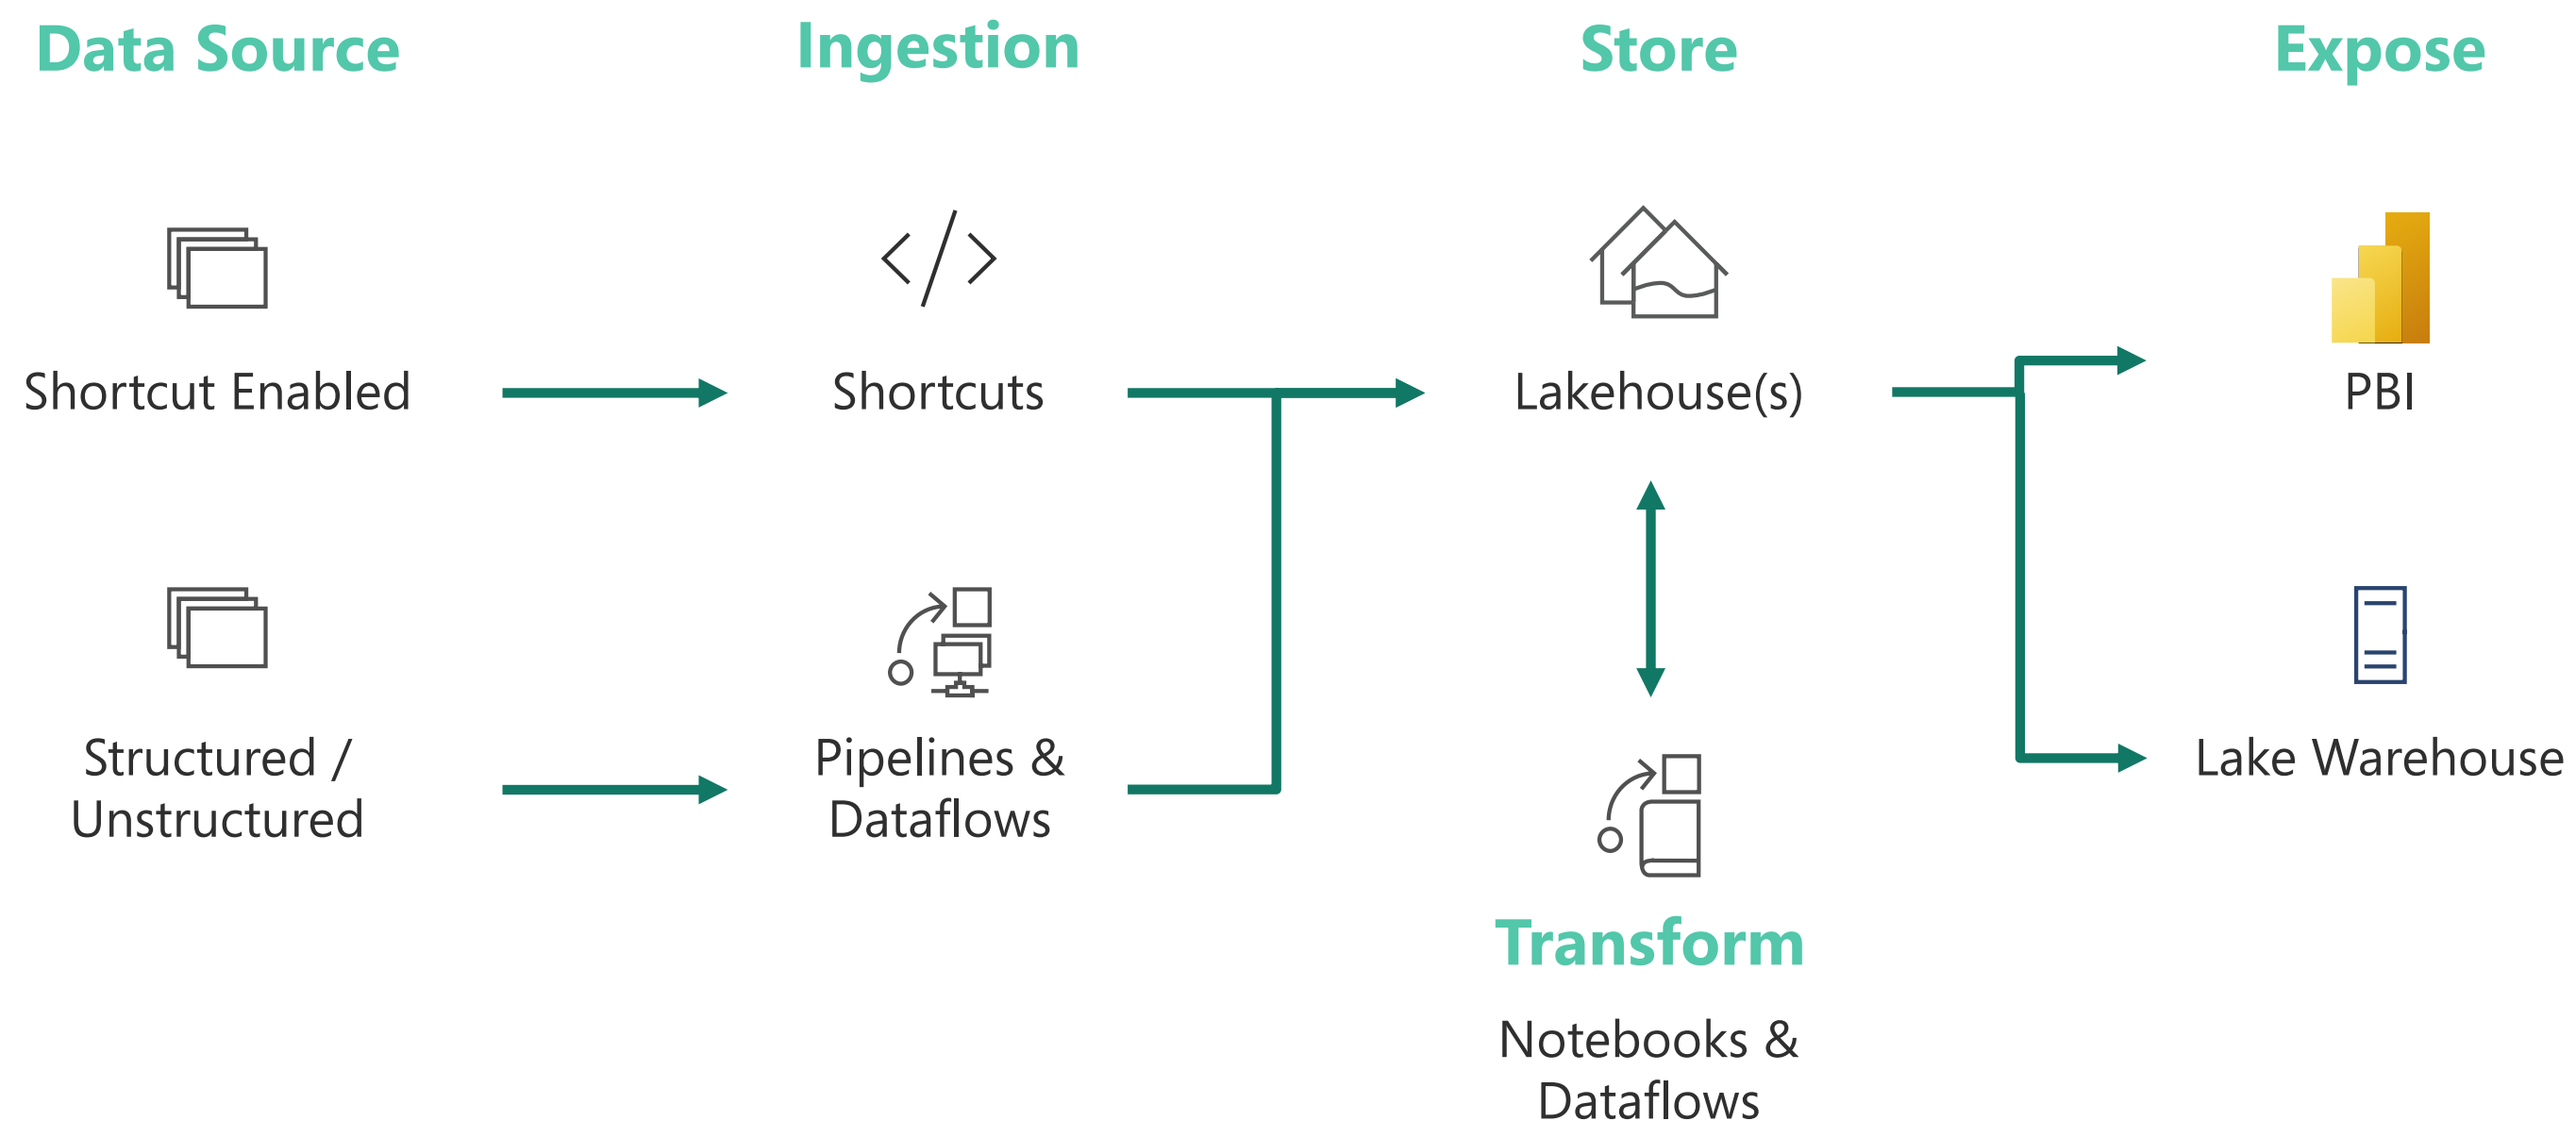
Please consider these scenarios as a guide to support your proof of concept. However, keep in mind that while these are the most common scenarios supported by Microsoft Fabric, there are many others that technical and business end-users can navigate through.

# 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 Direct Lake mode or SQL endpoints.

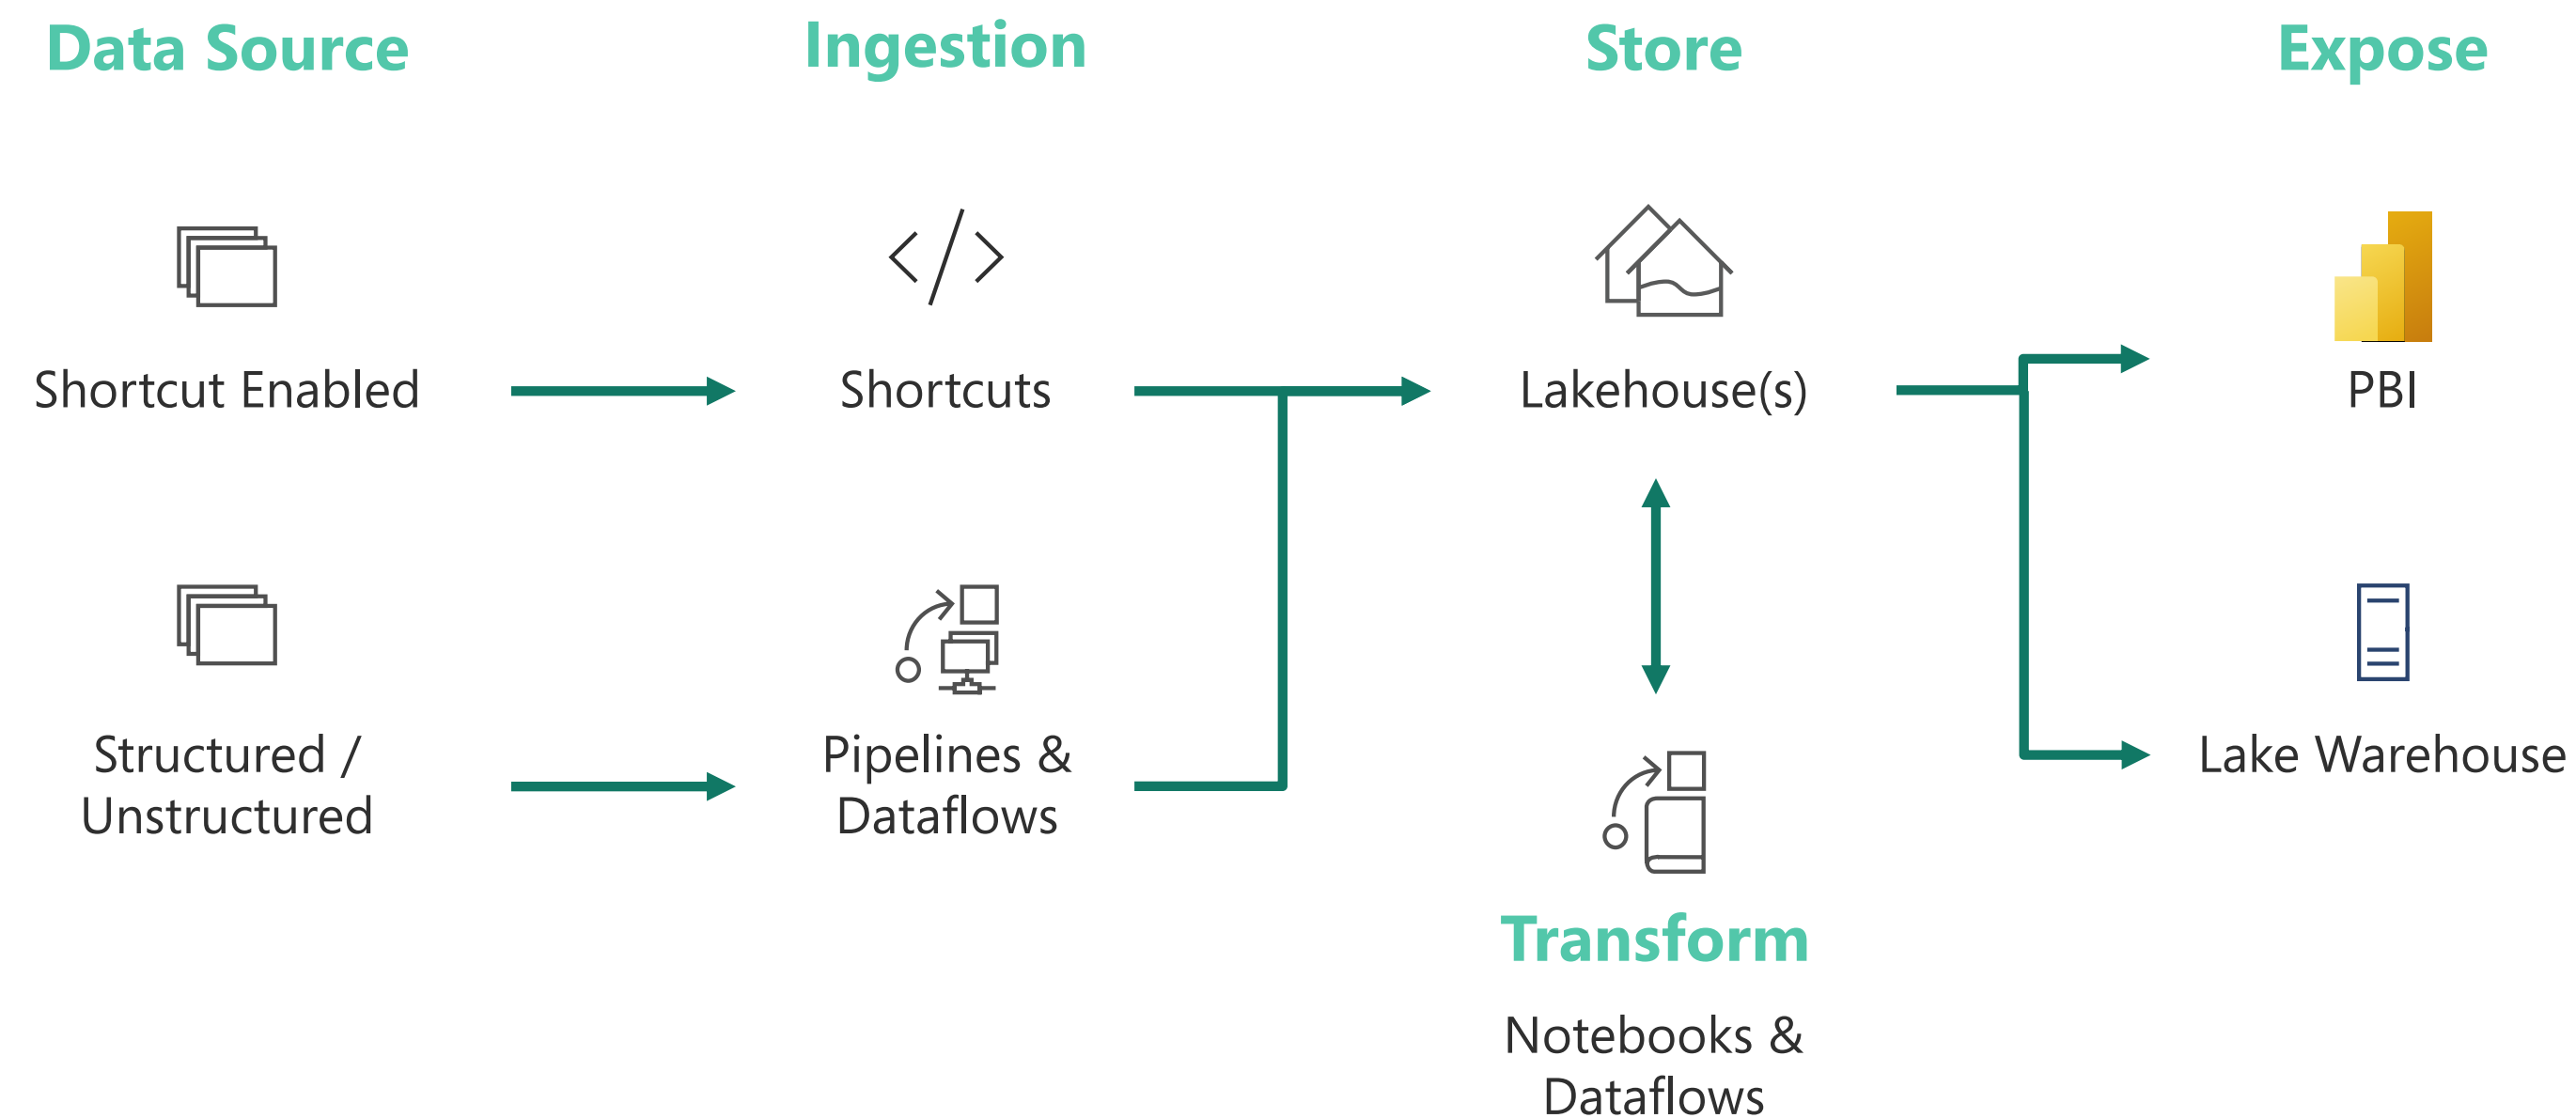


## Build and implement an end-to-end lakehouse for your organization:

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 Direct Lake – 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

# Lakehouse

## The lakehouse end-to-end architecture



**Data Sources** – Microsoft Fabric makes it easy and quick to connect to Azure Data Services, other cloud platforms, and on premises data sources to ingest data from.

**Ingestion** – With 200+ native connectors as part of the Microsoft Fabric pipeline and drag and drop data transformation with dataflow, you can quickly build insights for your organization. Shortcut is a new feature in Microsoft Fabric that provides a way to connect to existing data without having to copy or move it.

**Transform and Store** – Microsoft Fabric standardizes on Delta Lake format, that means all the engines of Microsoft Fabric can read and work on the same dataset stored in OneLake – no need for data duplicity. This storage allows you to build lakehouses using a medallion architecture or data mesh based on your organizational need. For transformation, you can choose either low-code or no-code experience with pipelines/dataflows or notebook/Spark for a code first experience.

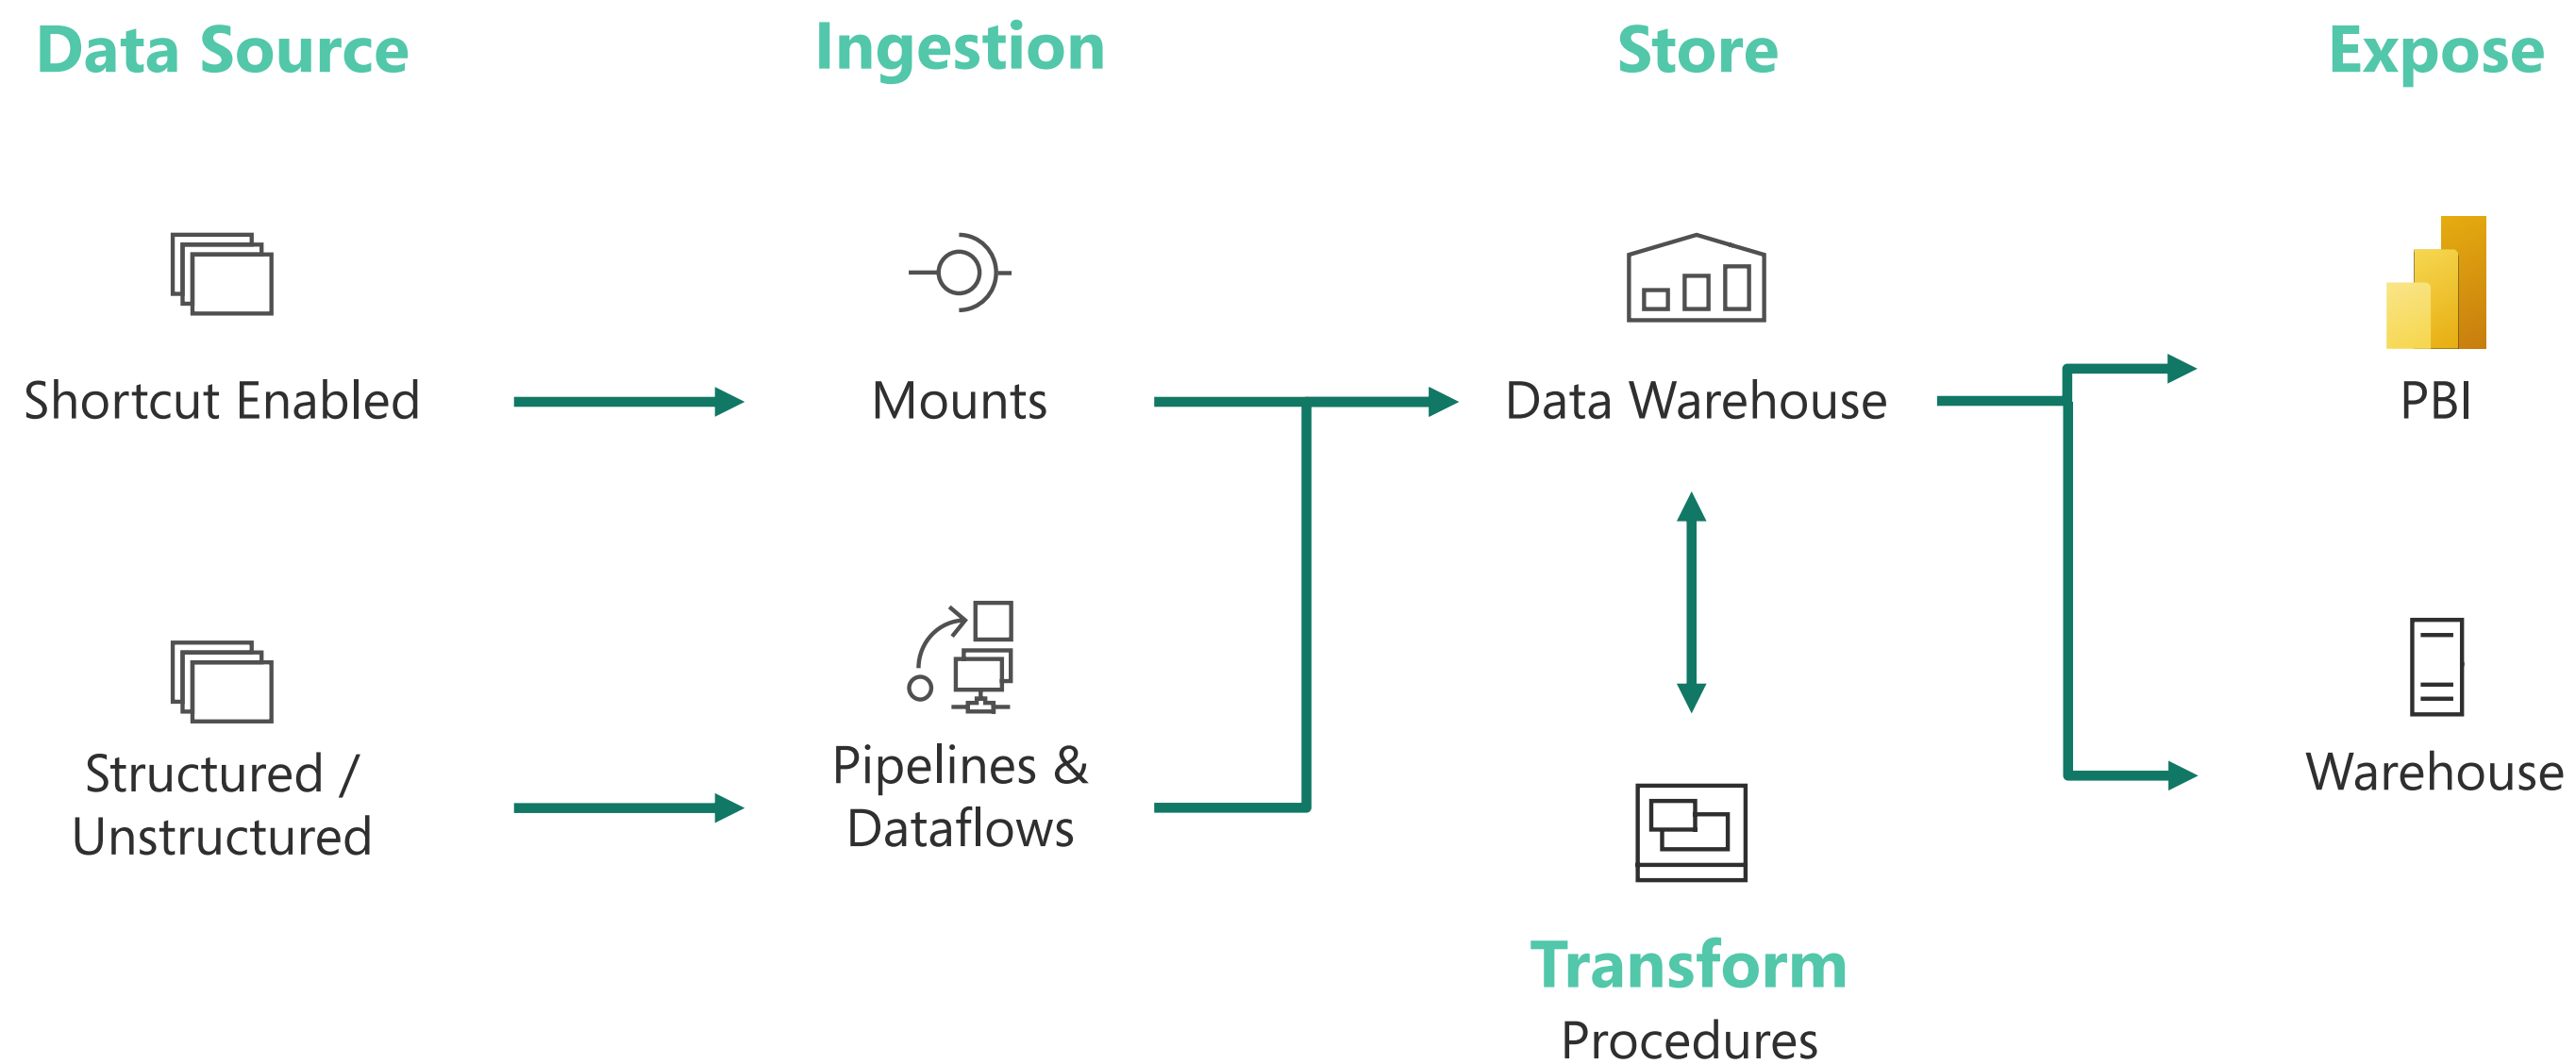
**Consume** – Data from your lakehouses can be consumed by Power BI, an industry leading business intelligence tool for reporting and visualization. Each Lakehouse comes with a built-in TDS/SQL endpoint for easily connecting to and querying data in the lakehouse tables from other reporting tools, when needed. When a Lakehouse is created a secondary item, called a Warehouse will be automatically generated at the same time with the same name as the Lakehouse and this Warehouse item provides you with the TDS/SQL endpoint.

# 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 Direct Lake mode or SQL endpoints.

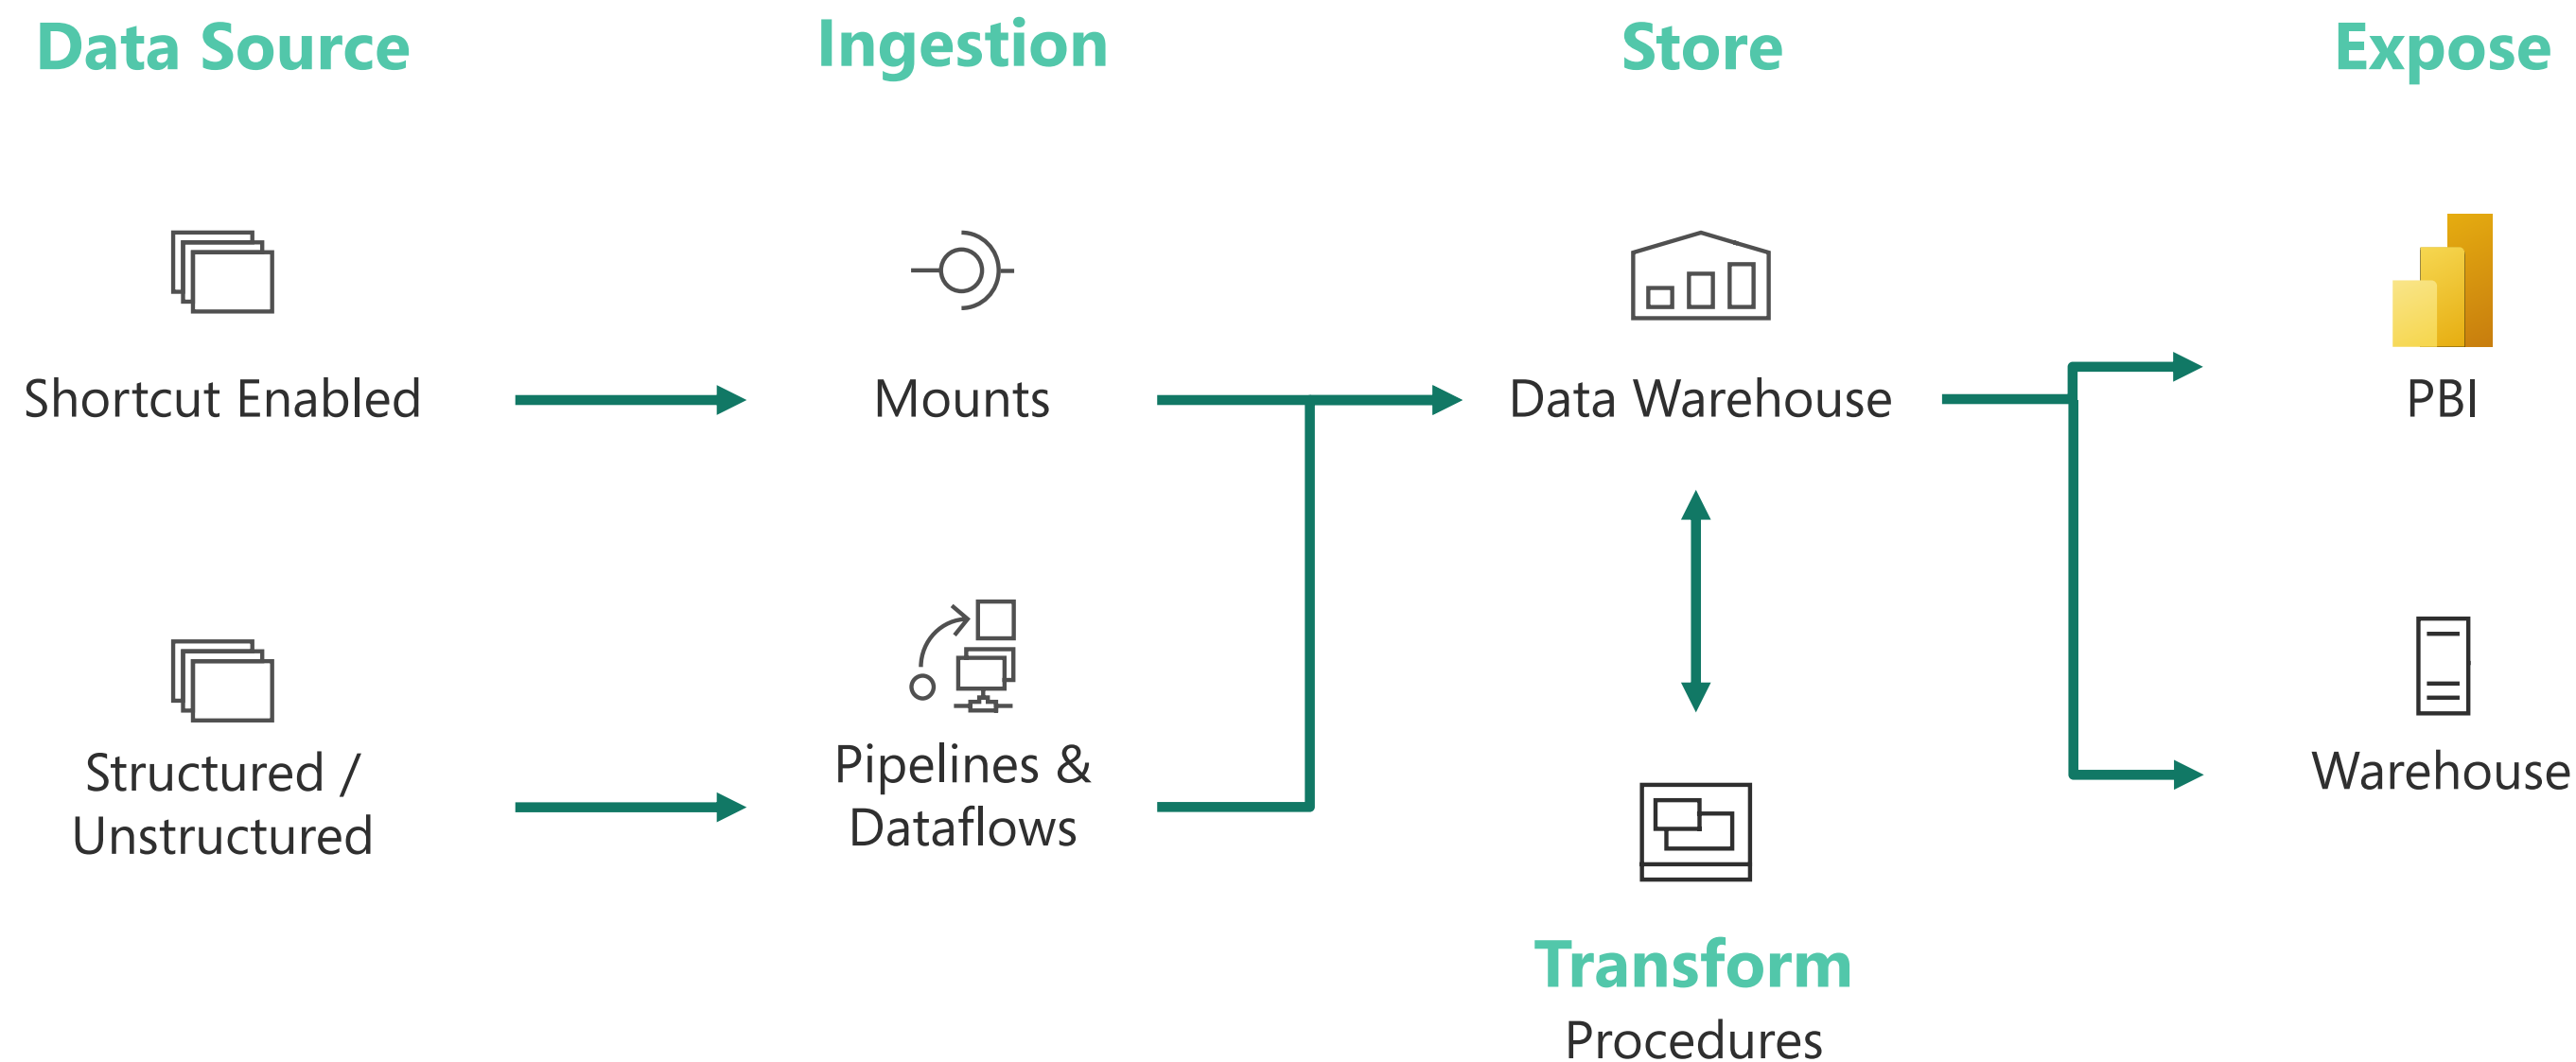


## Build and implement an end-to-end data warehouse for your organization:

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 Direct Lake mode to analyze the data in place
9. Cleanup resources by deleting the workspace and other items

# Data Warehouse

The data warehouse end-to-end architecture:



**Data Sources** – Microsoft Fabric makes it easy and quick to connect to Azure Data Services, other cloud platforms, and on premises data sources to ingest data from.

**Ingestion** – With 200+ native connectors as part of the Microsoft Fabric pipeline and drag and drop data transformation with dataflow, you can quickly build insights for your organization. Shortcut is a new feature in Microsoft Fabric that provides a way to connect to existing data without having to copy or move it.

**Transform and Store** – Microsoft Fabric standardizes on Delta Lake format, that means all the engines of Microsoft Fabric can read and work on the same dataset stored in OneLake – no need for data duplicity. This storage allows you to build a data warehouse or data mesh based on your organizational need. For transformation, you can choose either low-code or no code experience with pipelines/dataflows or use T-SQL for a code first experience.

**Consume** – Data from your lakehouses can be consumed by Power BI, an industry leading business intelligence tool for reporting and visualization. Each data warehouse comes with a built-in TDS/SQL endpoint for easily connecting to and querying data from other reporting tools, when needed. When a data warehouse is created a secondary item, called a default dataset, will be automatically generated at the same time with the same name of the data warehouse to start visualizing data with just a couple of mouse clicks.



# Data Science

The Data Science analytics scenario can be ingested similarly to the Lakehouse and Data Warehouse paths.

Once the data is ingested, it is cleaned and prepared using notebooks and then stored in the lakehouse with medallion structure.

After the data is cleaned and stored, machine learning models can be trained and tested directly on the lakehouse.

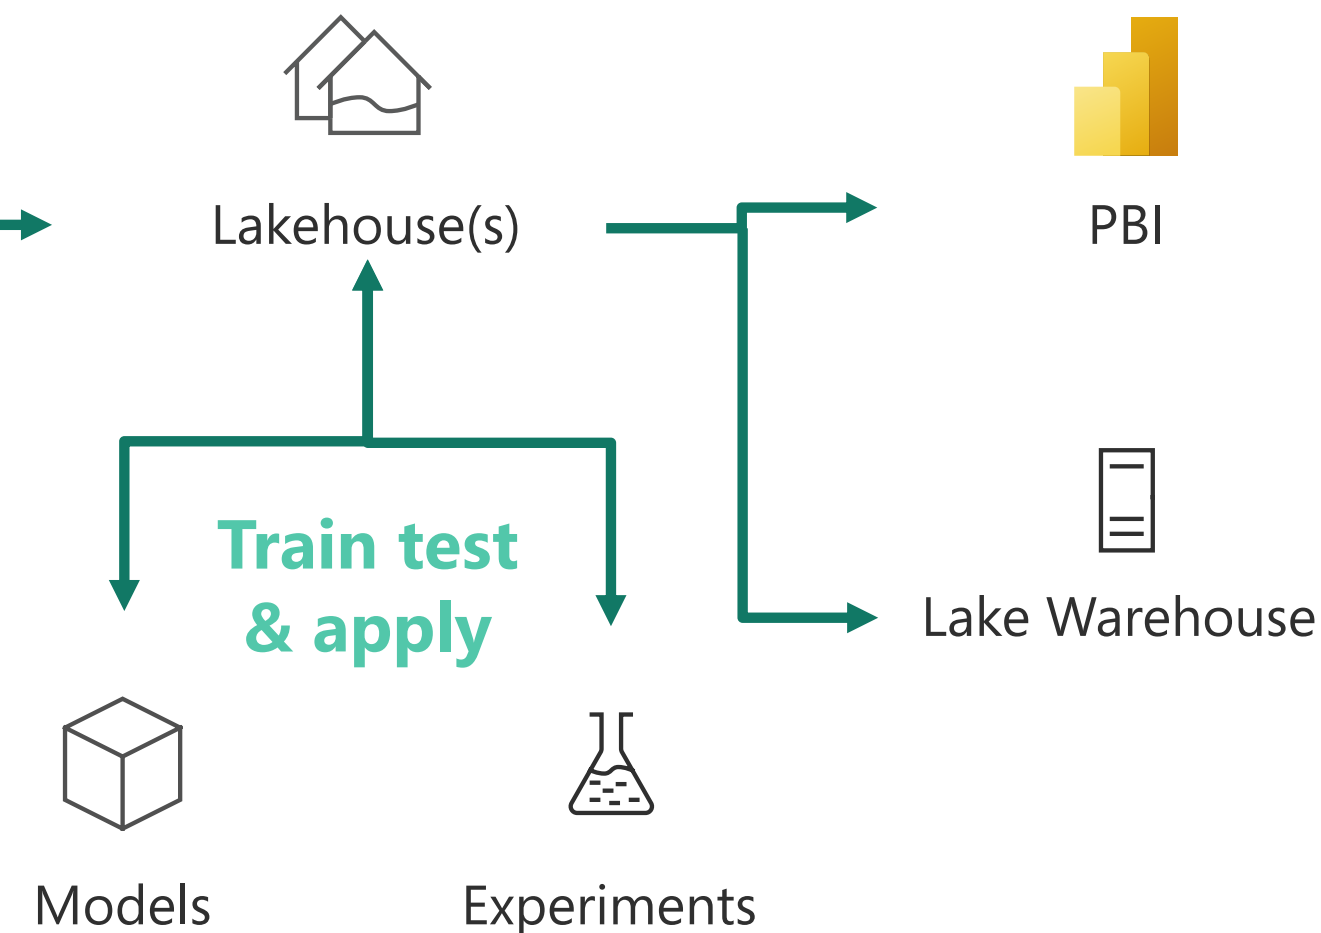
## Data Source



## Clean & prepare



## Store



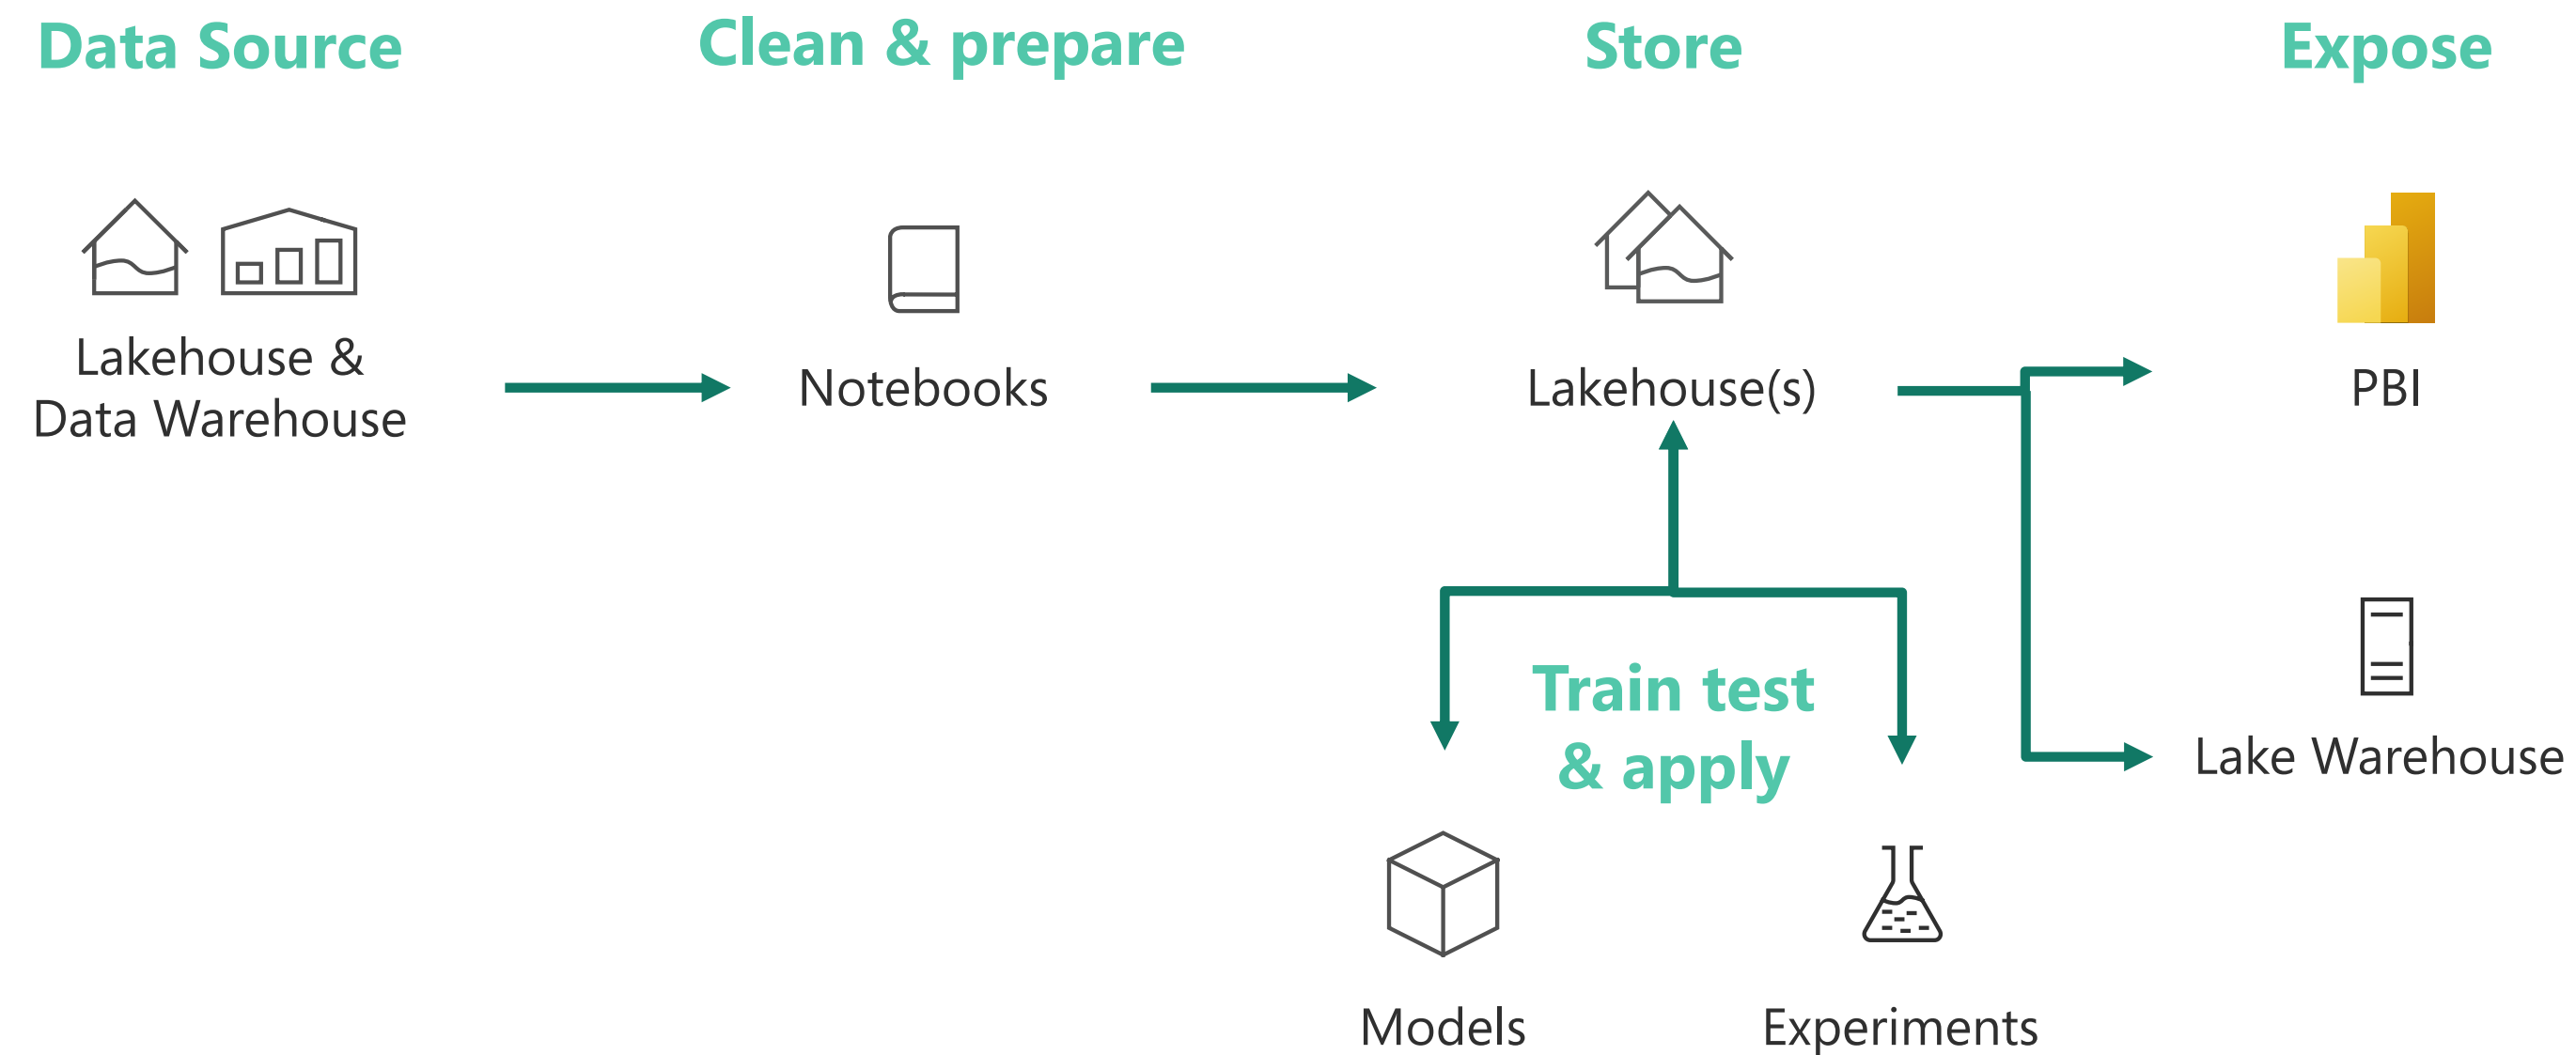
## Expose

**In this scenario, consider performing the following activities:**

1. Use the Microsoft Fabric notebooks for data science scenarios
2. Ingest data into Microsoft Fabric lakehouse using Apache Spark
3. Load existing data from the lakehouse delta tables
4. Clean and transform data using Apache Spark
5. Create experiments and runs to train a machine learning model
6. Register and track trained models using MLflow and the Microsoft Fabric UI
7. Run scoring at scale and save predictions and inference results to the lakehouse
8. Visualize predictions in Power BI using Direct Lake

# Data Science

## The data science end-to-end scenario



**Data Sources** – Microsoft Fabric makes it easy and quick to connect to Azure Data Services, other cloud platforms, and on-premises data sources to ingest data from. Using Microsoft Fabric Notebooks you can ingest data from the built-in Lakehouse, Data Warehouse, and Power BI Datasets as well as various Apache Spark and Python supported custom data sources.

**Explore, Clean & Prepare** – The Data Science experience on Microsoft Fabric supports data cleansing, transformation, exploration and featurization by leveraging built-in experiences on Spark as well as Python based tools like Data Wrangler and SemPy Library.

**Models & Experiments** – Microsoft Fabric enables you to train, evaluate and score machine learning models by using built-in Experiment and Model artifacts with seamless integration with MLflow for experiment tracking and model registration/deployment. Microsoft Fabric also features capabilities for model prediction at scale (PREDICT) to gain and share business insights.

**Storage** – Microsoft Fabric standardizes on Delta Lake, which means all the engines of Microsoft Fabric can interact with the same dataset stored in lakehouse. This storage layer allows you to store both structured and unstructured data that support both file-based storage as well as tabular format. The datasets and files stored can be easily accessed via all Microsoft Fabric workload artifacts like notebooks and pipelines.

**Expose Analysis and Insights** – Data from the lakehouse can be consumed by Power BI, an industry leading business intelligence tool for reporting and visualization. Data persisted in the lakehouse can also be visualized in notebooks using Spark or Python native visualization libraries like matplotlib, seaborn, plotly, and more. Data can also be visualized using SemPy library that supports built-in, task-specific visualizations for the semantic data model, dependencies and their violations, and classification and regression use cases.

# Real Time Analytics

Unlike the Data Science, Lakehouse, and Data Warehouse analytics scenarios, streaming data can be ingested into the Microsoft Fabric in several ways to achieve real time analytics.

Users can leverage Event Hub, IoT Hub, pipelines, dataflows, notebooks, or open-source products like Kafka, Logstash, and more.

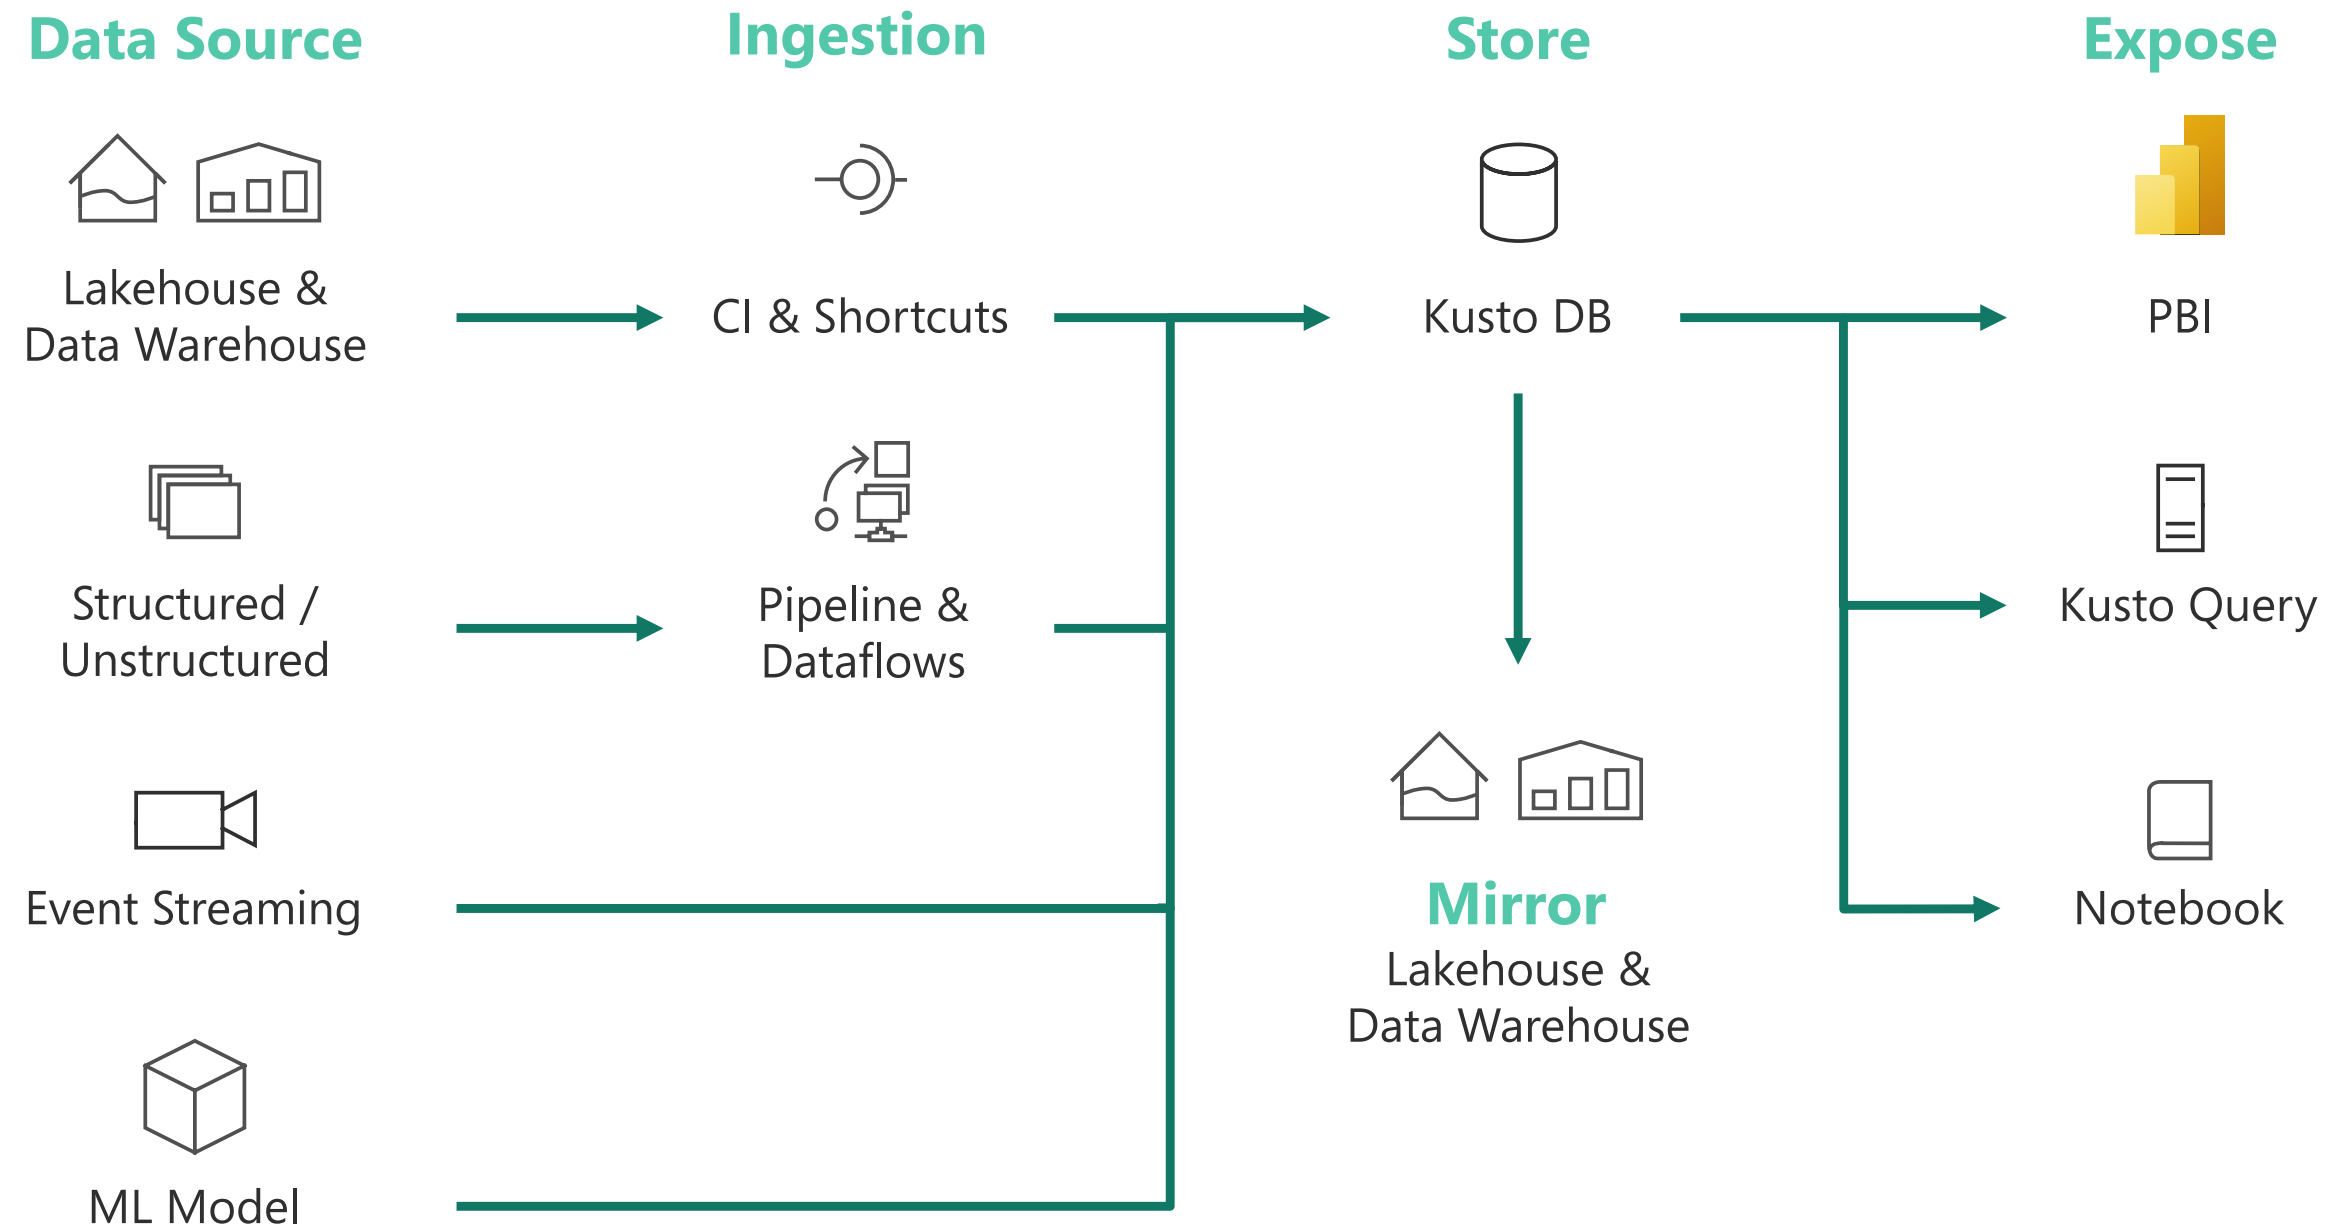
Once ingested into Microsoft Fabric, streaming data can be stored in Kusto DB and mirrored into the lakehouse. After the data has been stored, machine learning models can be trained and tested directly on the lakehouse with experiments.

Like the other scenarios, business users can analyze and visualize the data with Power BI using the Direct Lake mode or SQL endpoints. Data can also be exposed through KQL or notebooks using Spark.

**In this scenario, consider performing the following activities:**

1. Create a KQL Database
2. Create Eventstream
3. Stream data from Eventstream to KQL Database
4. Check your data with sample queries
5. Save queries as a KQL Queryset
6. Create a Power BI report
7. Create a OneLake shortcut

## The end-to-end Real-Time Analytics scenario




### PHASE 3

# Build & Scale

Delivering accelerated data value is not an overnight process. It involves significant organizational alignment. It necessitates building a solid and reliable data foundation. And lastly it requires consistently delivering business value and enabling decision making to build trust through incremental success over time.

# Building a foundation of alignment

We recommend the following steps to ensure the foundation for your journey is solid and you can help your entire organization embrace the culture of insight and innovation.



## 1. Get aligned with your organization's leaders and data stewards

- Build confidence with department leaders
- Collaborate on your business value roadmap
- Develop your infrastructure roadmap

## 2. Map your current and future data estate

- Document the details of your current data estate
- Ideate on planned transformation and outcomes

## 3. Establish a network of transformation champions

- Assemble a team of cross-functional subject matter experts

# Aligning with your organization's leaders and data stewards



Alignment with your organization's department leaders is crucial to a successful, iterative and value-driven data estate transformation. You'll likely need to start by helping these important stakeholders see the broader vision of what is possible and the economic impact associated with your aspirations.

According to the Harvard Business Review, data shows that "companies that embrace a data-driven culture experience a **4x improvement in revenue performance and better customer satisfaction**". The impact of a data culture on an organization is enormous.

Source: [How to lead a data-driven digital transformation by Harvard Business Review](#)

## Stakeholders crucial to data estate transformation


- Chief Financial Officer and Other Finance Leaders
- Chief Sales Officer and Other Sales Leaders
- Chief Marketing Officer and Other Marketing Leaders
- Human Resources Leaders
- Chief Operations Officer and Other Operations Leaders (including Supply Chain, Ecommerce, Manufacturing, and more)
- Customer Service Leaders
- Research and Development Leaders
- Corporate Strategy and Planning Leaders



# Conversations and questions to drive your business value roadmap

We recommend running ideation sessions with department leaders to identify the initiatives that can drive the most business value once your data estate is unified. Consider an initiative that can drive some initial momentum and can be iterated on.

1. How can we enable innovation with data and analytics for your department?
2. What innovation can we enable?
3. How can I make your line of business more efficient?
4. What immediate initiatives can we collaborate on and how will we measure the impact?
5. What business value will we deliver over time? Cost reduction? Revenue generation?
6. How can I better enable your line of business to use data in decision-making?
7. What security and governance concerns should we be aware of?
8. Is there third-party or IoT data that could enrich your insights?



## WORKSHEET

### Create a list of your initiatives to drive business value.

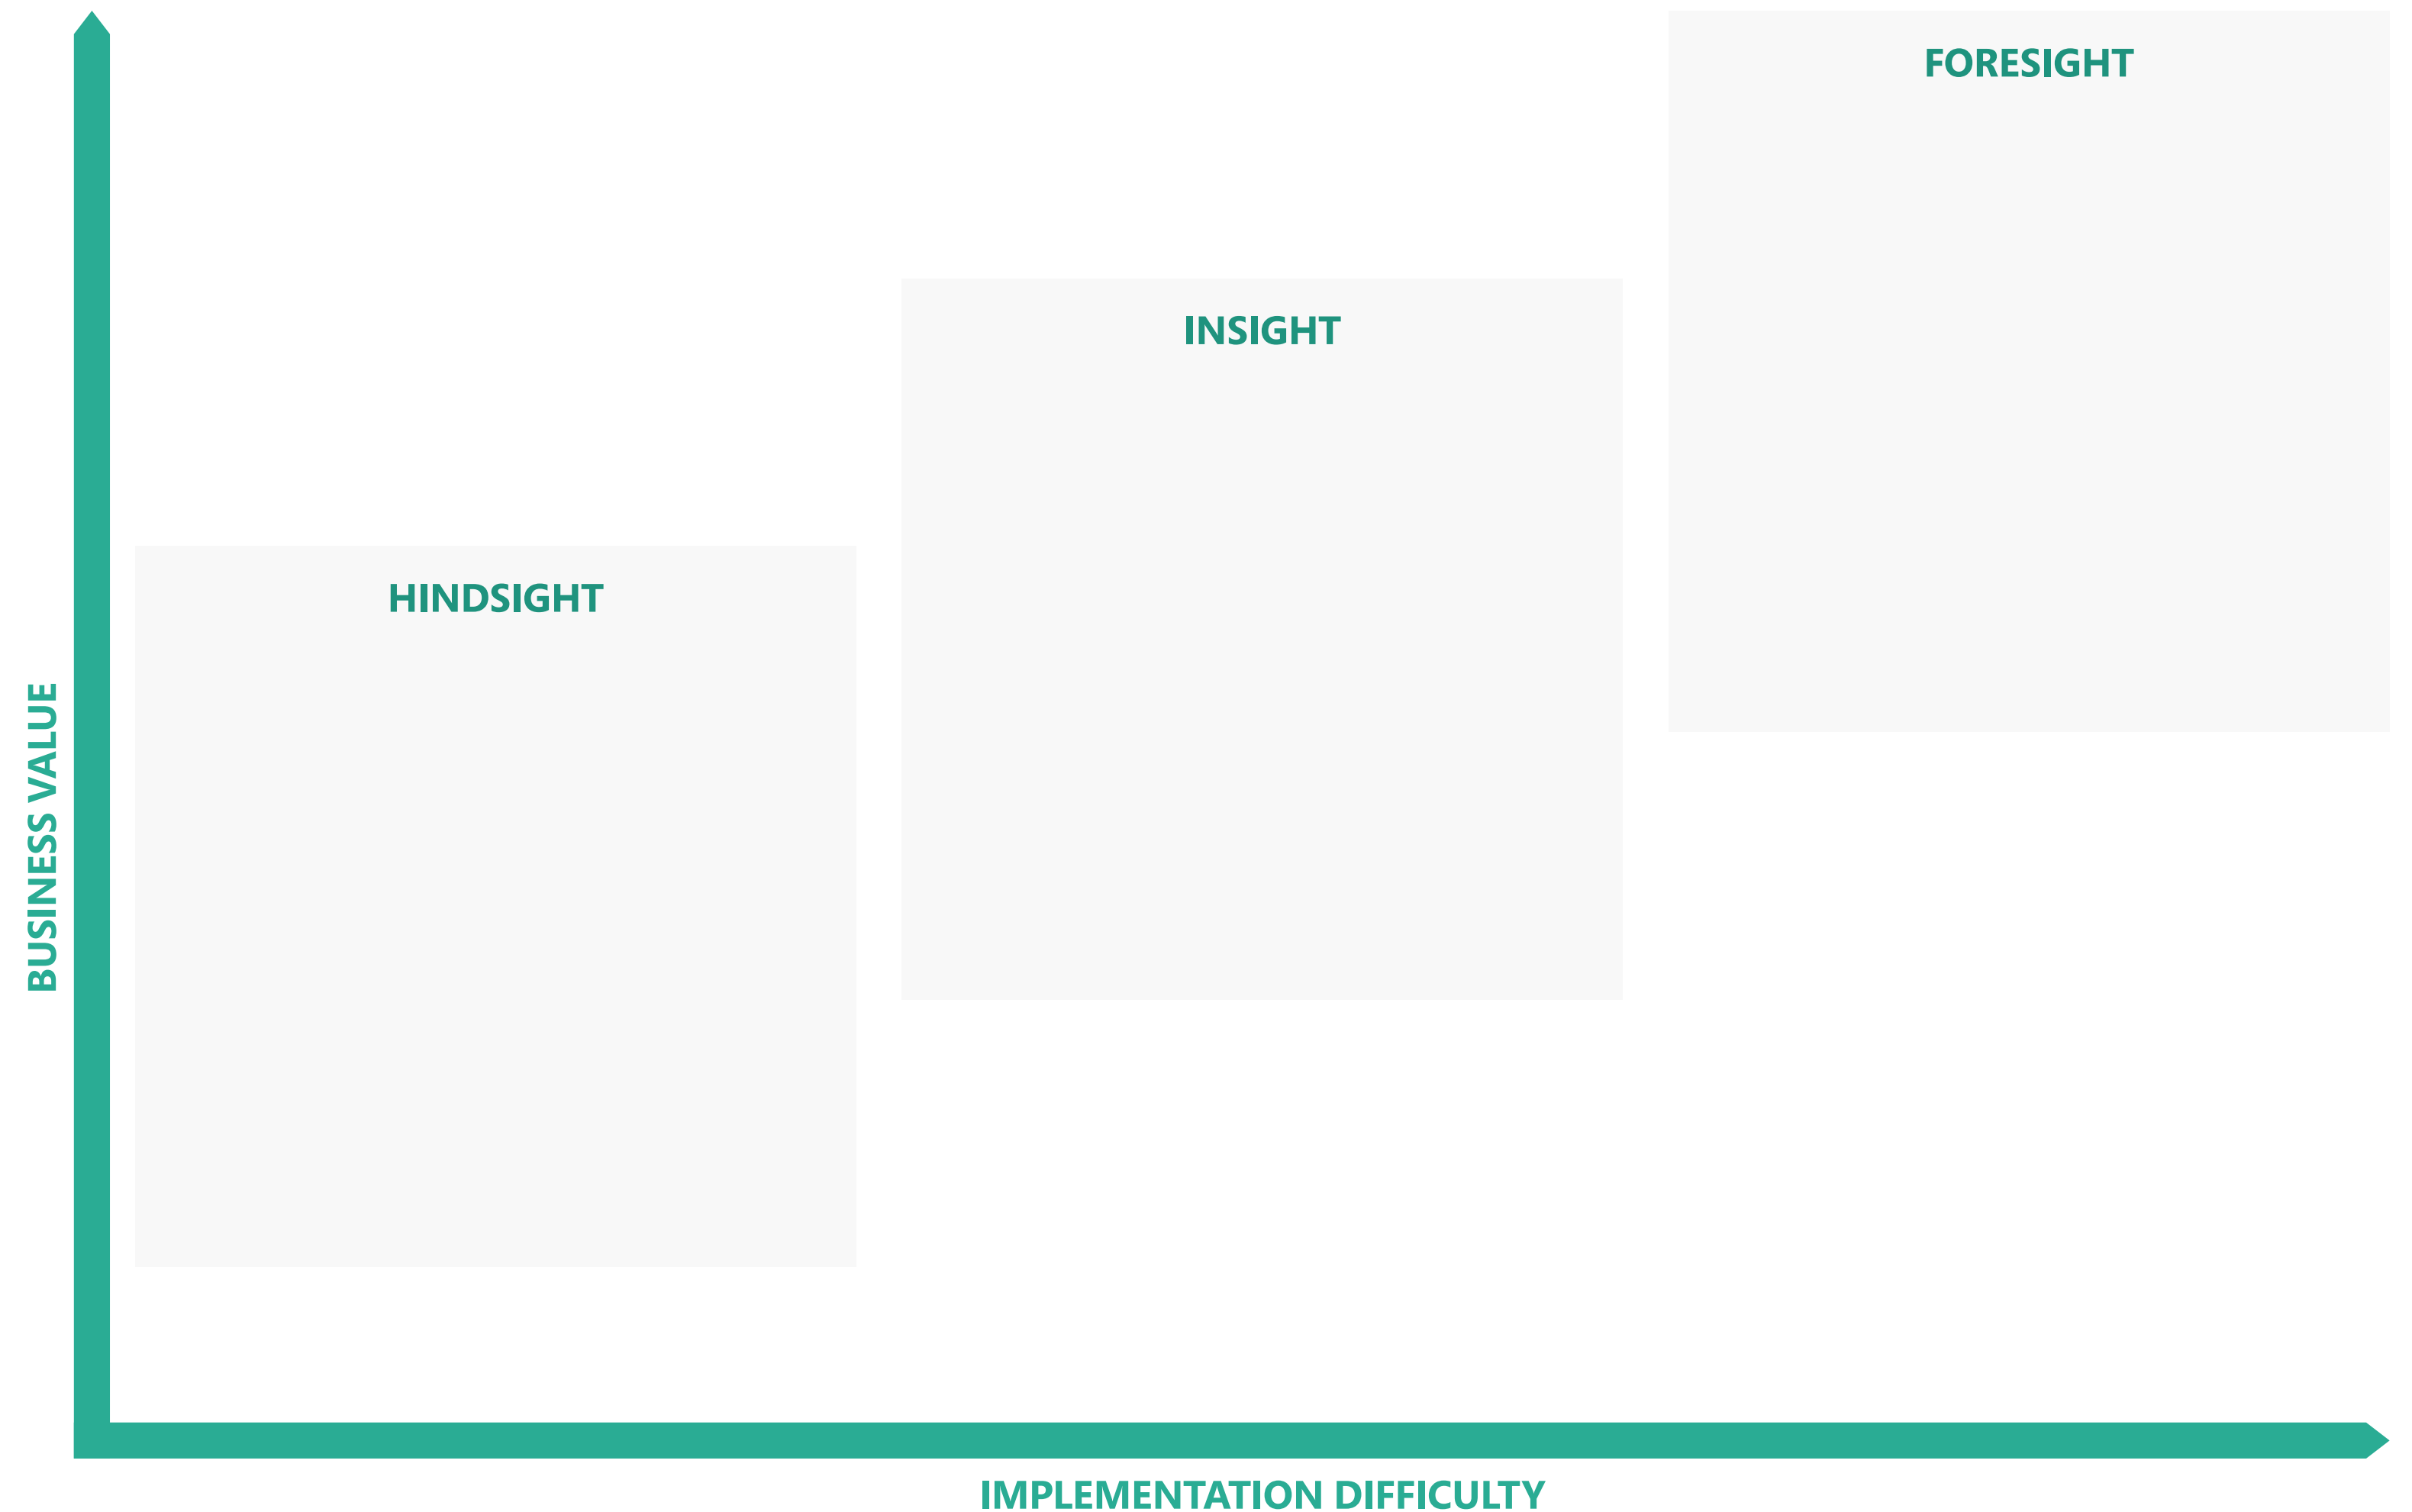
Categorize each initiative by the type of analytics it will support in the future.

**Hindsight:** is both descriptive and diagnostic, answering what happened and why it happened


**Insight:** is predictive, to consider what might happen

**Foresight:** is prescriptive, to suggest a course of action

This will be the start of your backlog and will help you prioritize based on the impact it has the potential to make. While you will ultimately settle on one or two ideas to begin with, once you have started and begun seeing success, this backlog of ideas will form a roadmap of what to do next.







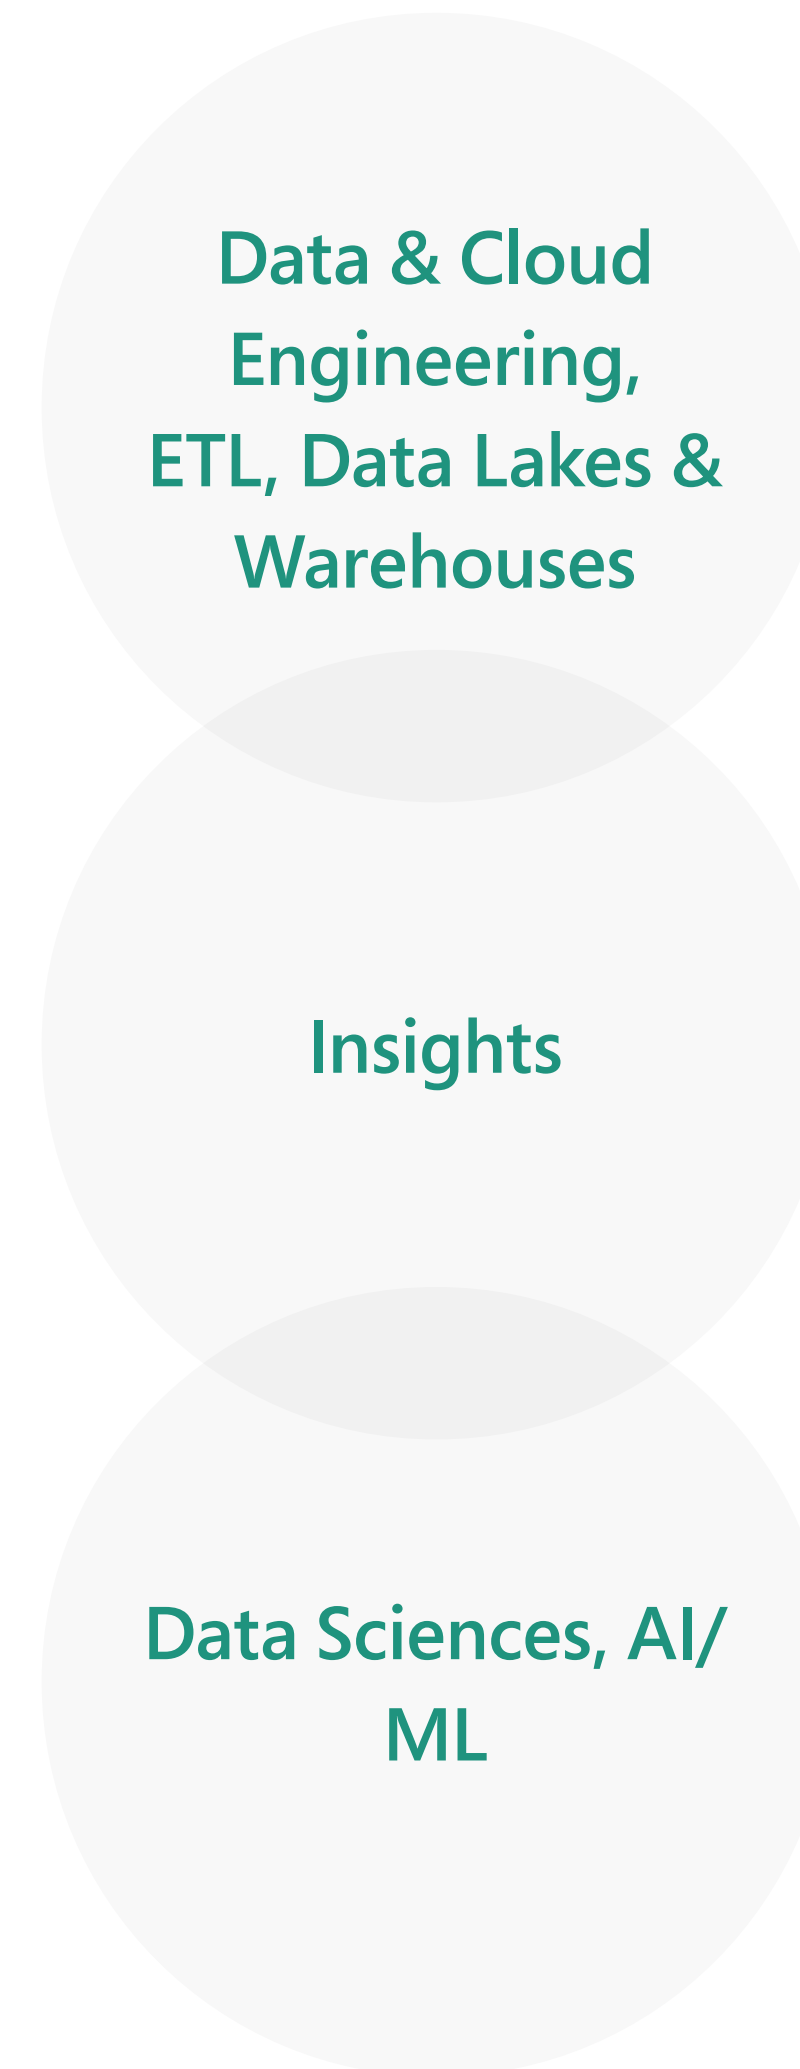
## WORKSHEET


# Conversations and considerations to drive your infrastructure roadmap

Alignment between the Chief Data Officer, Chief Technology Officer, and Chief Information Officer is imperative in the transformation of your data estate. Ensuring partnership through clear functional alignment as well as decision-making authority during these initiatives will pave the way for a smooth journey.

### Key Considerations and Discussion Points for Alignment

- Budget Ownership
- Decision Making R&R
- Governance
- Compliance
- Quality Assurance
- Operating Models
- Processes
- Monetization





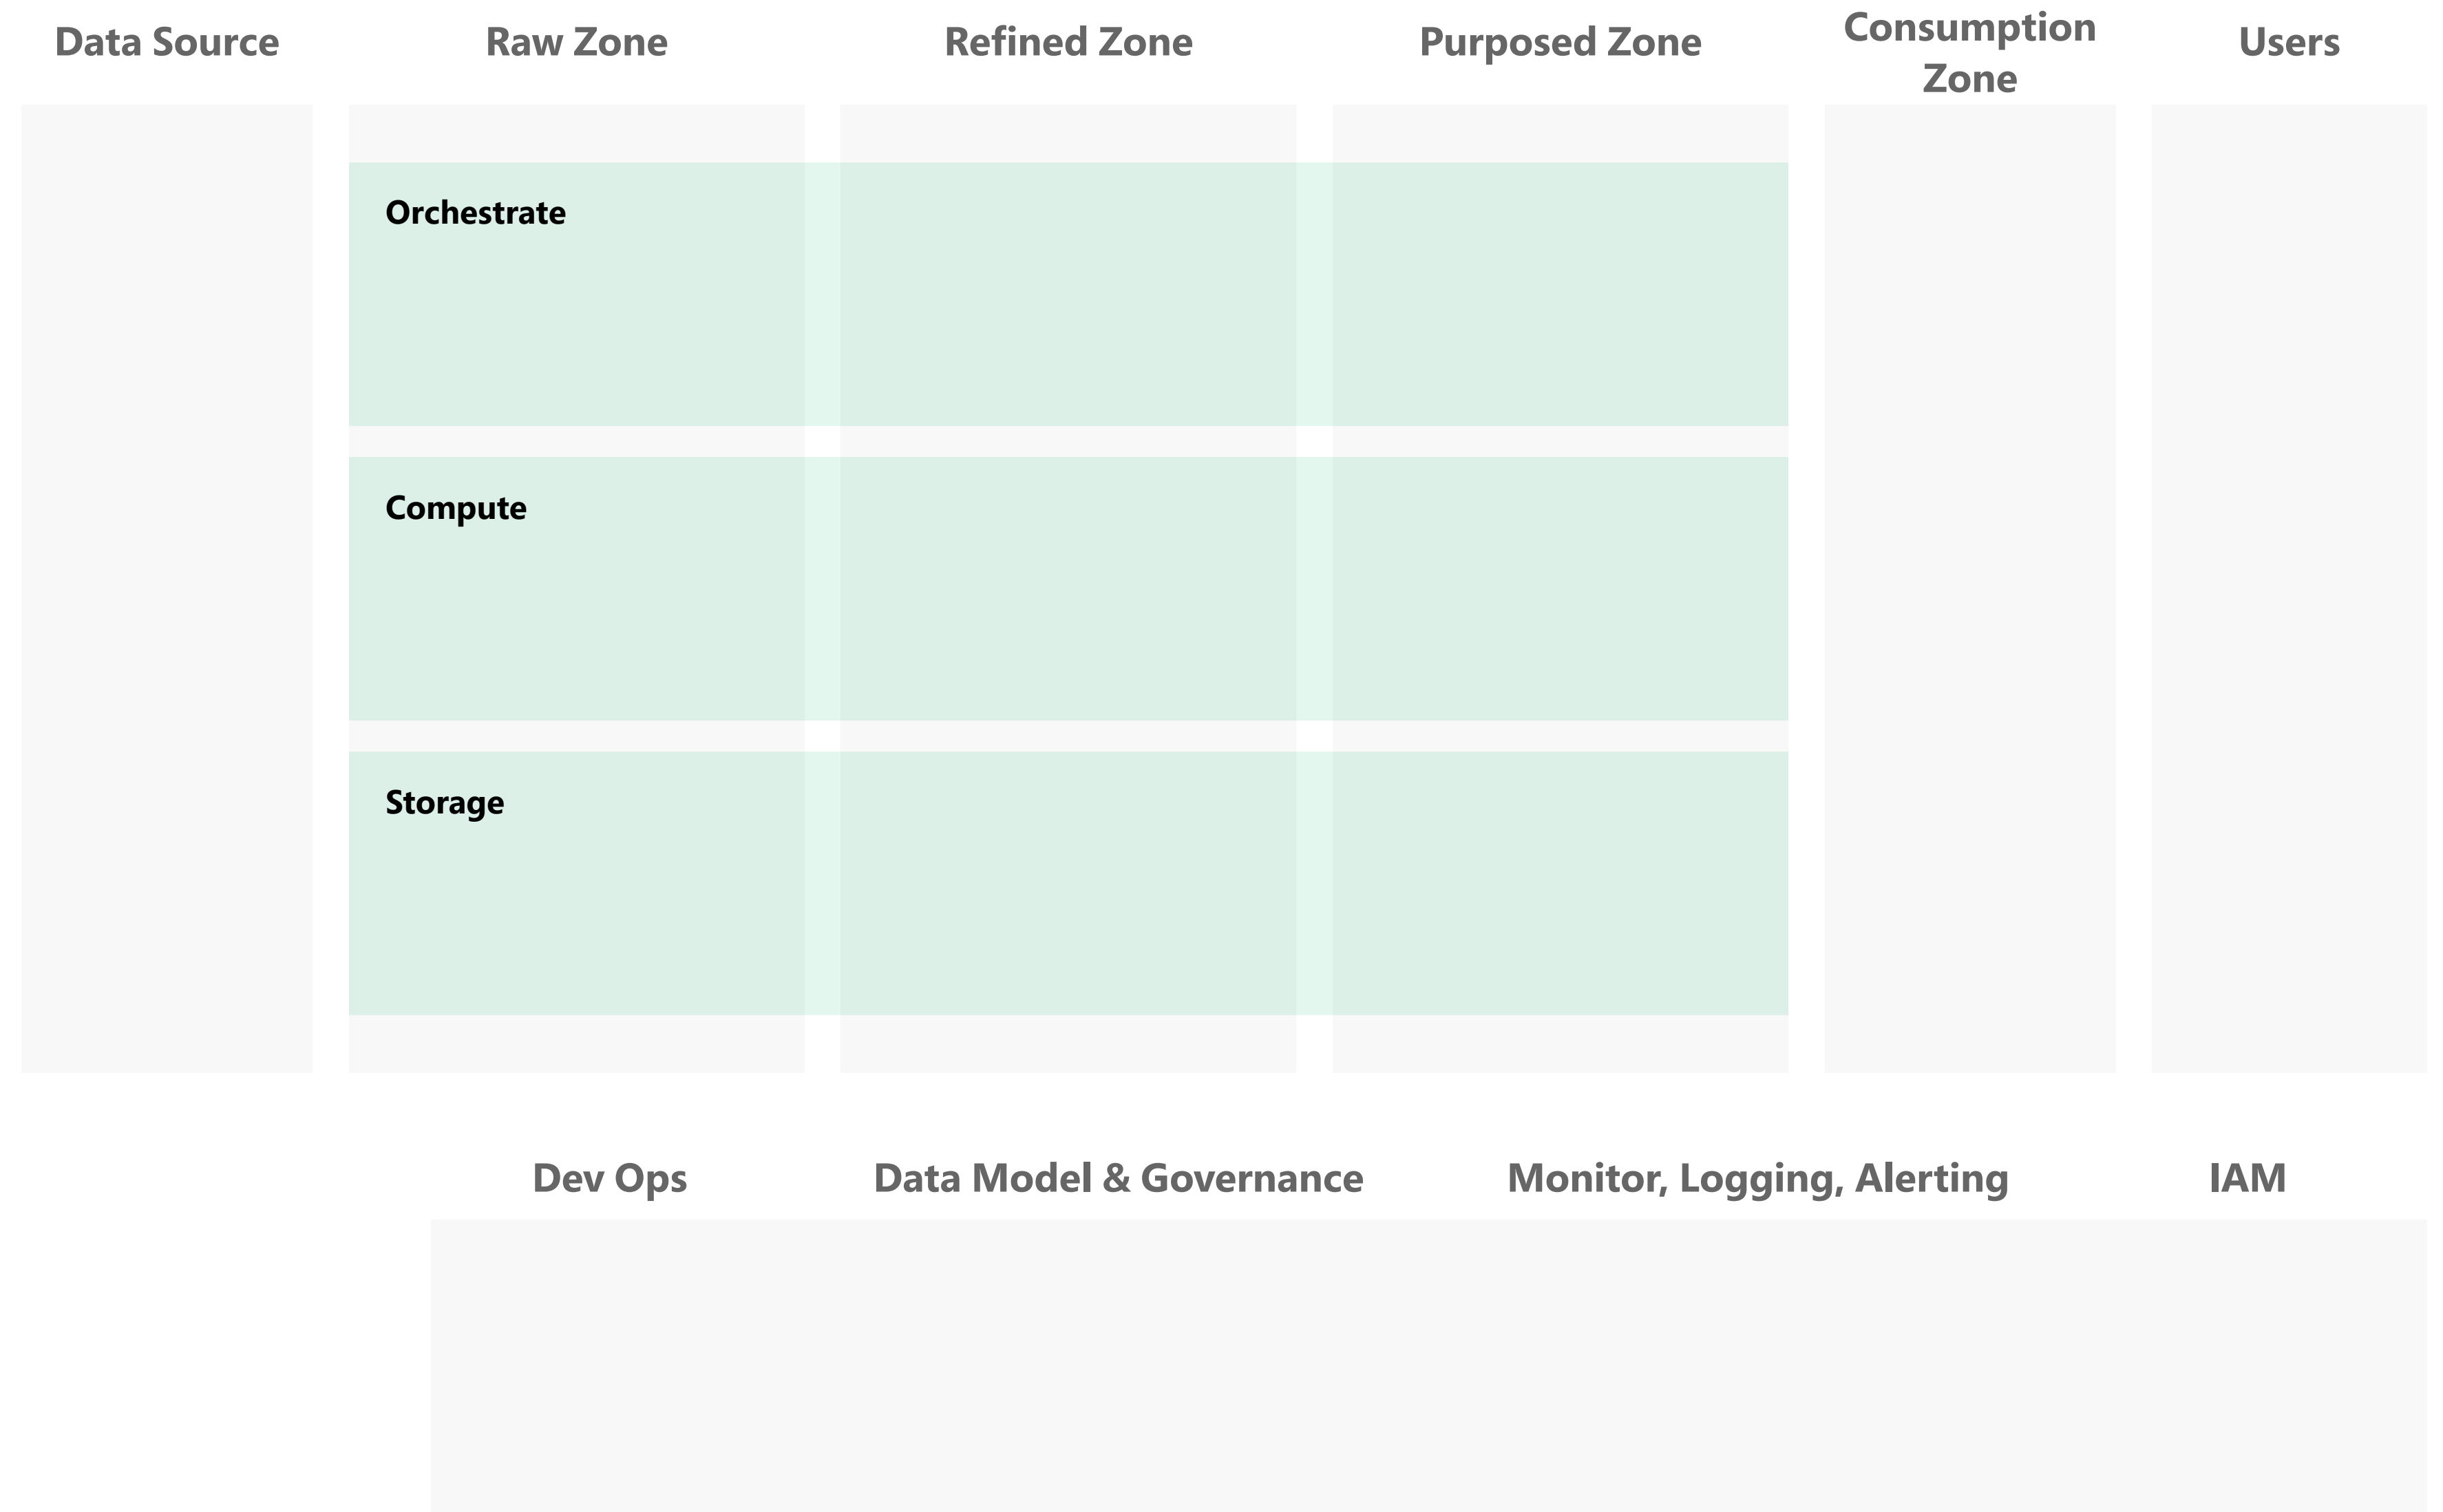
## WORKSHEET


# Mapping your current data and analytics estate

With your infrastructure team, begin documenting your current data estate. Be sure to capture which data sources, technologies, ELT/ETL scenarios, data pipelines, and appliances exist. Also, document how data is consumed by users.

### Document the details of your data estate.

Which data sources, technologies, ELT/ETL scenarios exist? How is data consumed by users?






**WORKSHEET**

# Ideate on planned transformation and outcomes

Begin to document the details of your data estate transformation. Use the inputs from your business value roadmap to ensure you've prioritized these initiatives in your phased infrastructure evolution.

Technology	In Scope/Out of Scope	Planned Transformation	Expected Outcomes



## WORKSHEET

# Establishing a network of transformation champions

Begin assembling your team of champions for your data estate transformation. Select people who will work well together, but also come from different roles and competencies in the organization and with various skills and experiences. You'll want to find a balance between functional expertise.

Team Member	Role	Functional Skillset



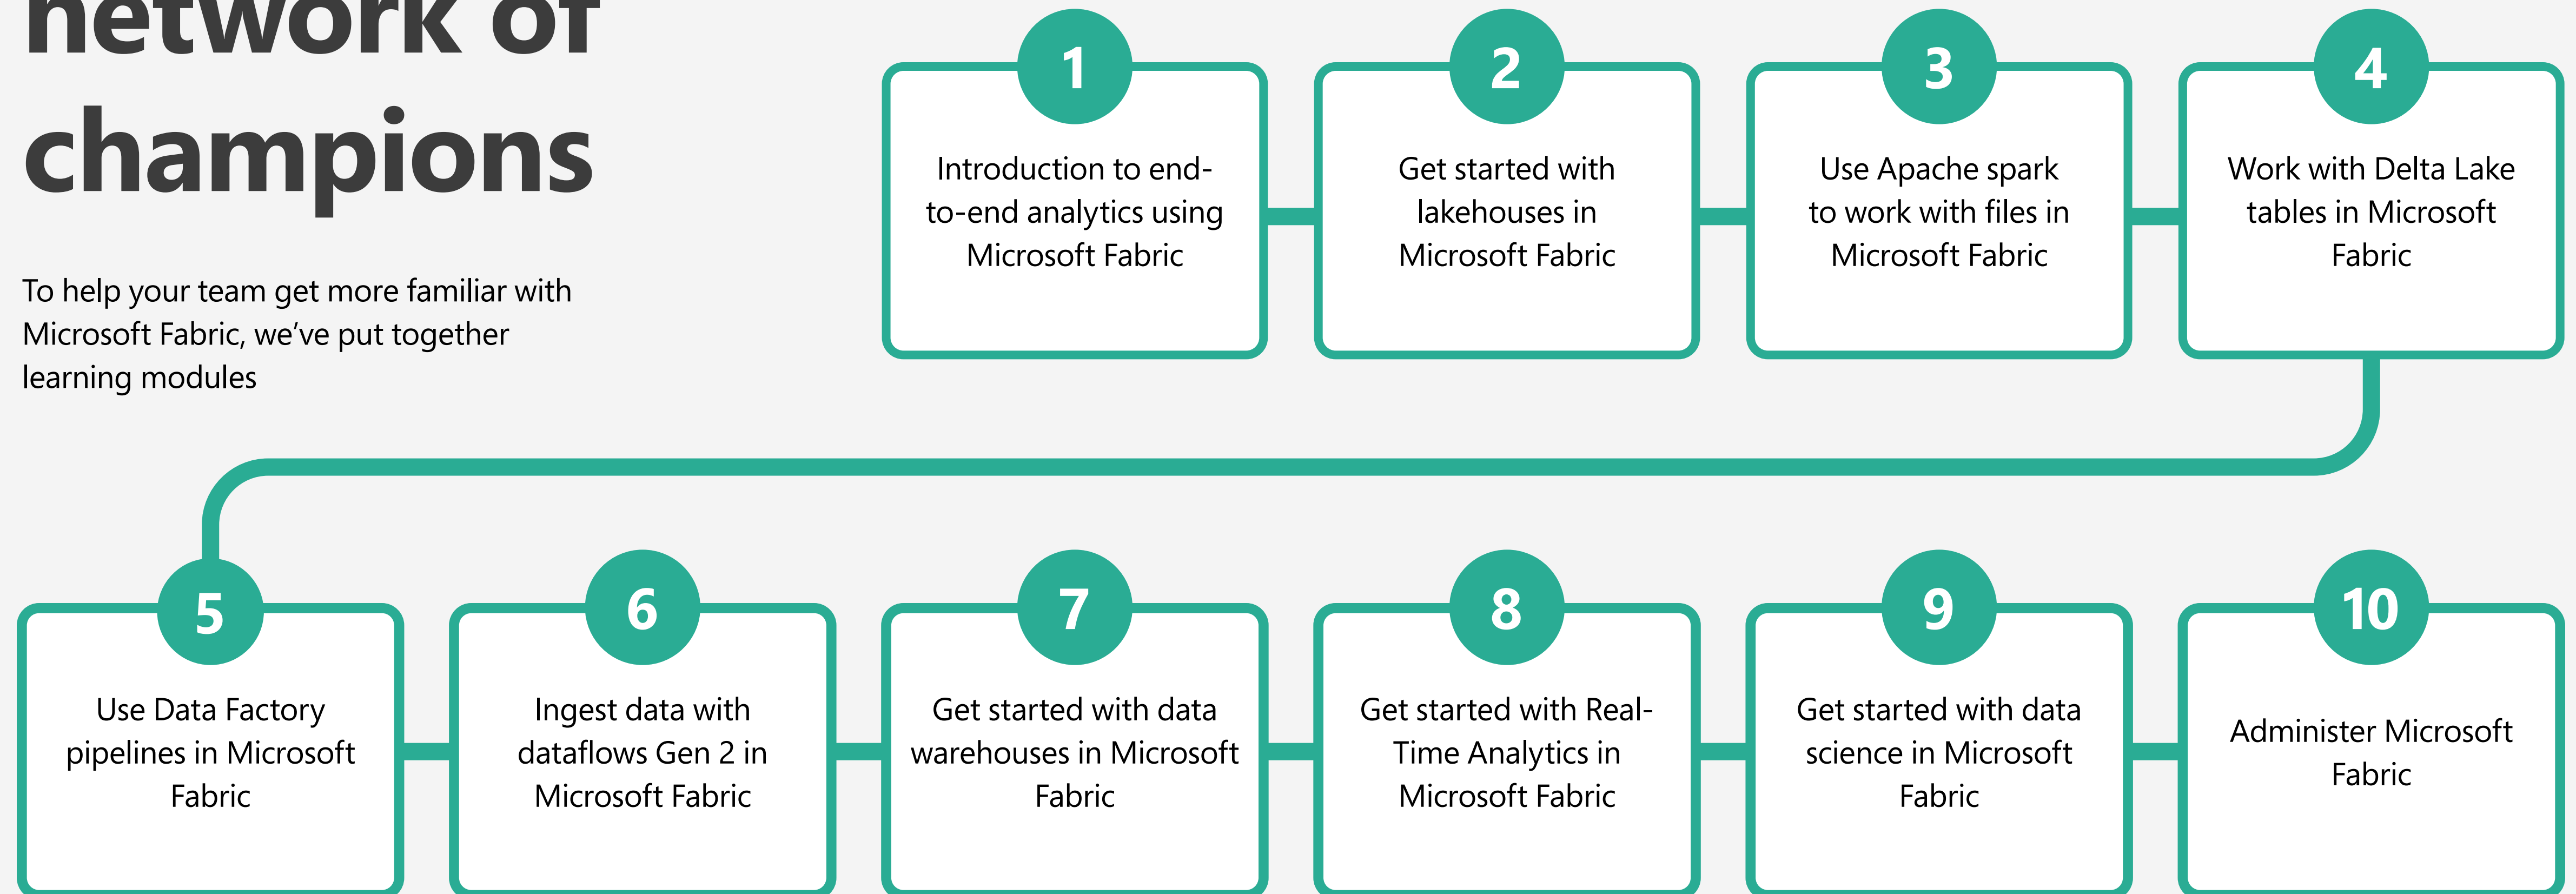
## We recommend the following roles as part of your innovation team.

- **Department leaders (VPs / Directors):** who can bring significant domain expertise
- **IT leaders:** Infrastructure admin and developers
- **Data practitioners:** Data scientists, data engineers, business analysts, data architects who will bring specialized knowledge of a process, function, or task
- **Data stewards:** Data professionals with knowledge of security, compliance, governance, and access management
- **Business users and data analysts:** Power BI power users



# Training your network of champions

To help your team get more familiar with Microsoft Fabric, we've put together learning modules



# You are on your way to transformative data value

Congratulations on following the journey to data estate transformation. You are well on your way to leading your organization to accelerated data value and innovation. As you continue your path, you may find these additional resources helpful.

Explore the Microsoft Fabric free trial today

Engage with Microsoft Fabric Community

Sign up for updates

Watch the digital event

Watch the webinar series