



2026
DYNAMICS
CON

Vegas



Microsoft Fabric


Data Analytics Platform



Stephanie Liang

Data Engineering Manager,
Alltech, Inc.

19.80	20.20	1.000
19.30	18.50	1.000
12.00	19.50	1.000



Power BI



Zach Flynn

BI & Analytics
Development Manager
Alltech, Inc.



Agenda



Alltech Data Brain

Medallion Architecture

Incremental Processing Strategy

Metadata Driven Framework


Modular & Scalable Design

DevOps & CI/CD

Security & Governance

Progress

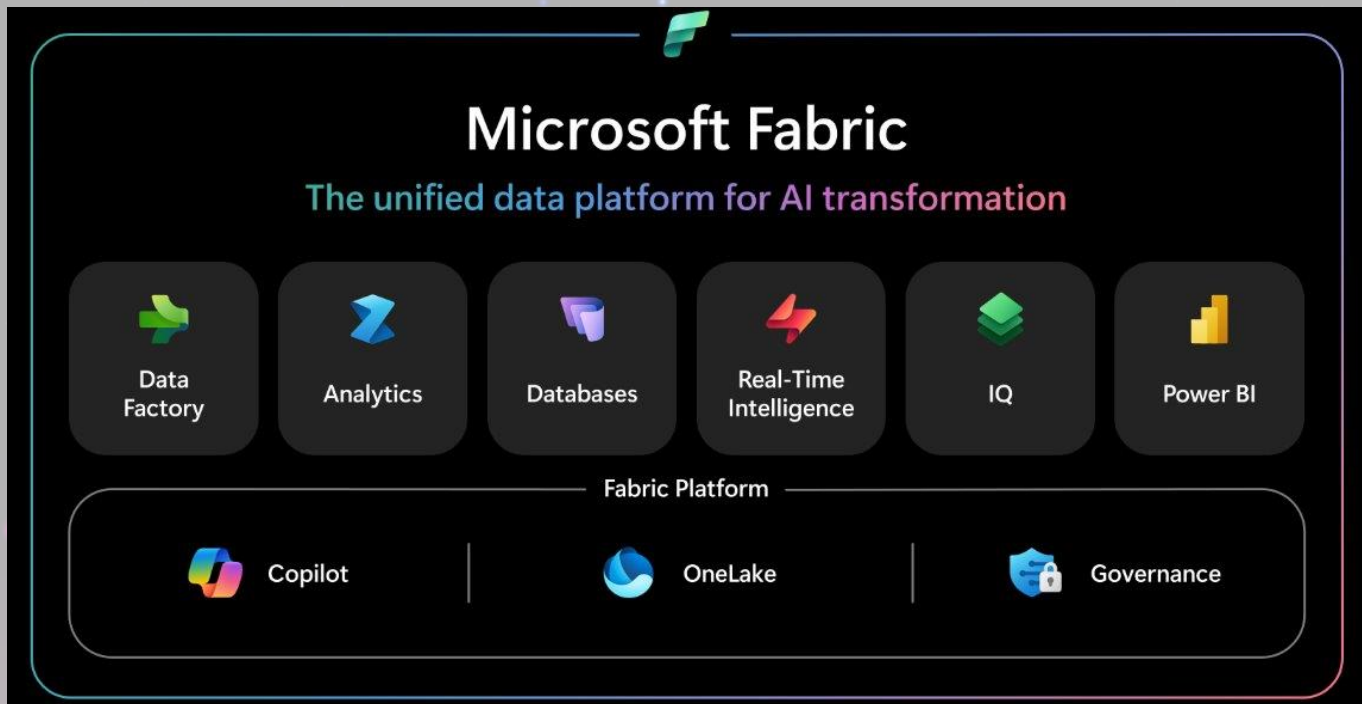
Building a Data Brain



Next Generation Data Platform - Microsoft Fabric

Microsoft Fabric

Microsoft Fabric is an **end-to-end analytics platform** unifying data ingestion, transformation, and reporting in a single SaaS experience.



Unified Data Platform

Combines data engineering, Data Factory, Data warehousing and reporting

OneLake

Centralized Data Lake

- ⇒ New ERP data: **GL, sub-ledger, AR/AP** (D365/AX)
- ⇒ Existing data: Sales, Purchasing, Inventory, 37+ Alltech platforms
- ⇒ External data: commodity, weather, customer info, market trends

Data-Driven Insights

Enables **self-service** analytics with less IT dependency

Fabric Terminology

OneLake



OneLake is a single, unified, logical data lake for the whole organization. Like **OneDrive**, OneLake is designed to be the single place for all analytics data.

Lakehouse



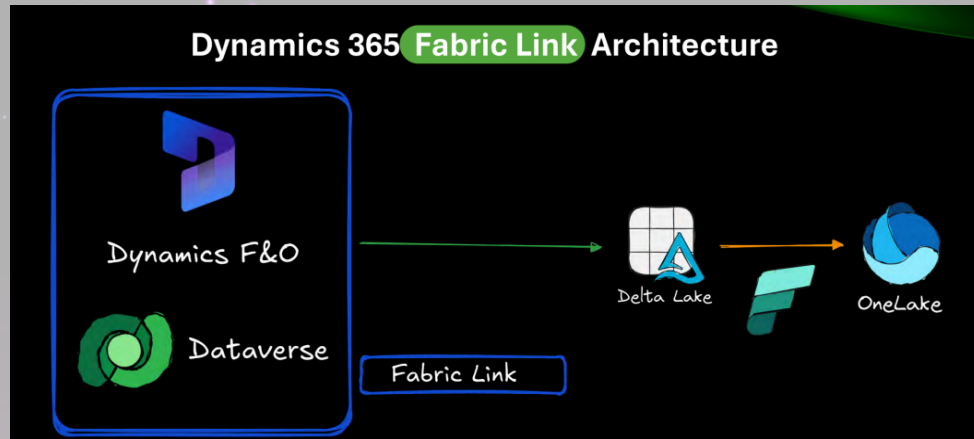
VS

Warehouse



They are within OneLake Platform

Fabric Link



Shortcut



A pointer to the other storage, enable users to access data without copying data.

Mirroring



It's like data looking into a mirror – It keeps a live and always up to date copy of your data in Fabric.

Notebook



How data is processed

Notebook is where you can write and run code, it supports: PySpark, Python, Spark SQL, Scala, SparkR and T-SQL

Pipeline



When and in what order it runs

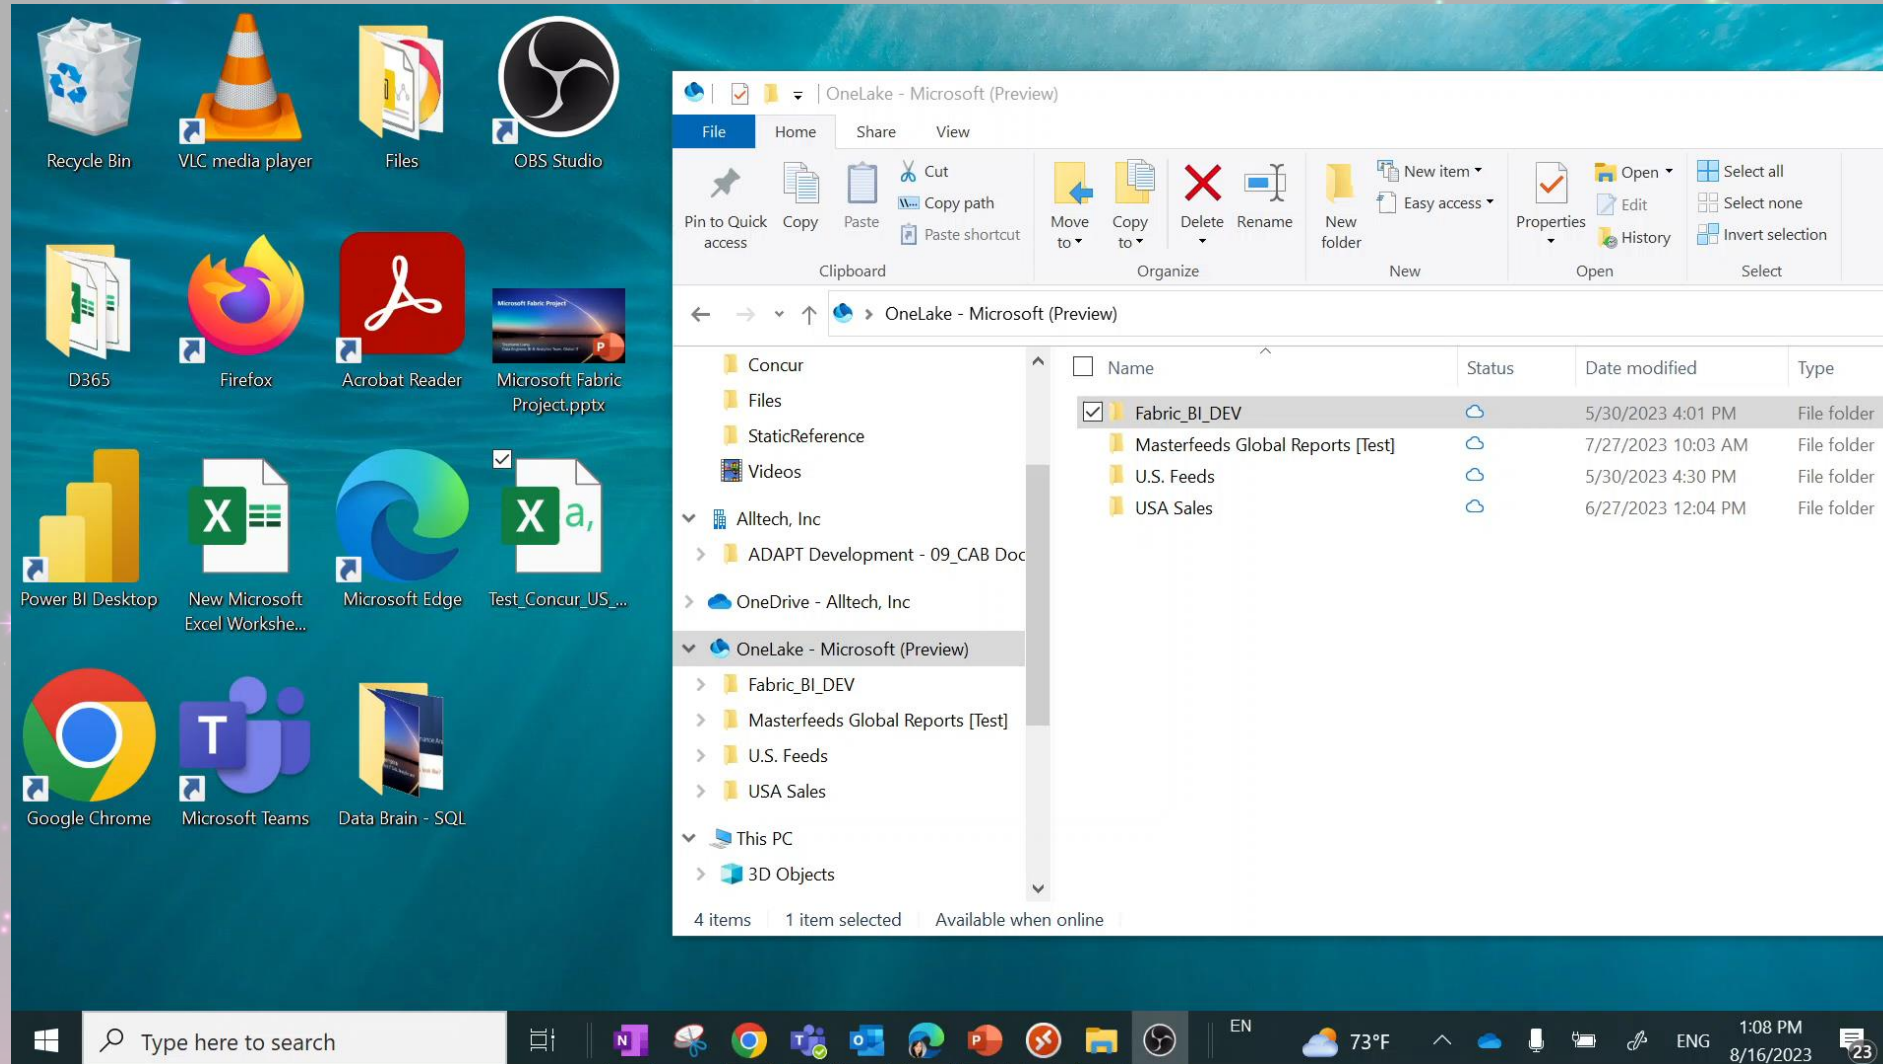
We use it to orchestrates process. It runs automatically based on a schedule.

Artifact

Artifacts refer to various objects such as dataflows, datasets, reports, and pipelines that are created and managed within the platform.

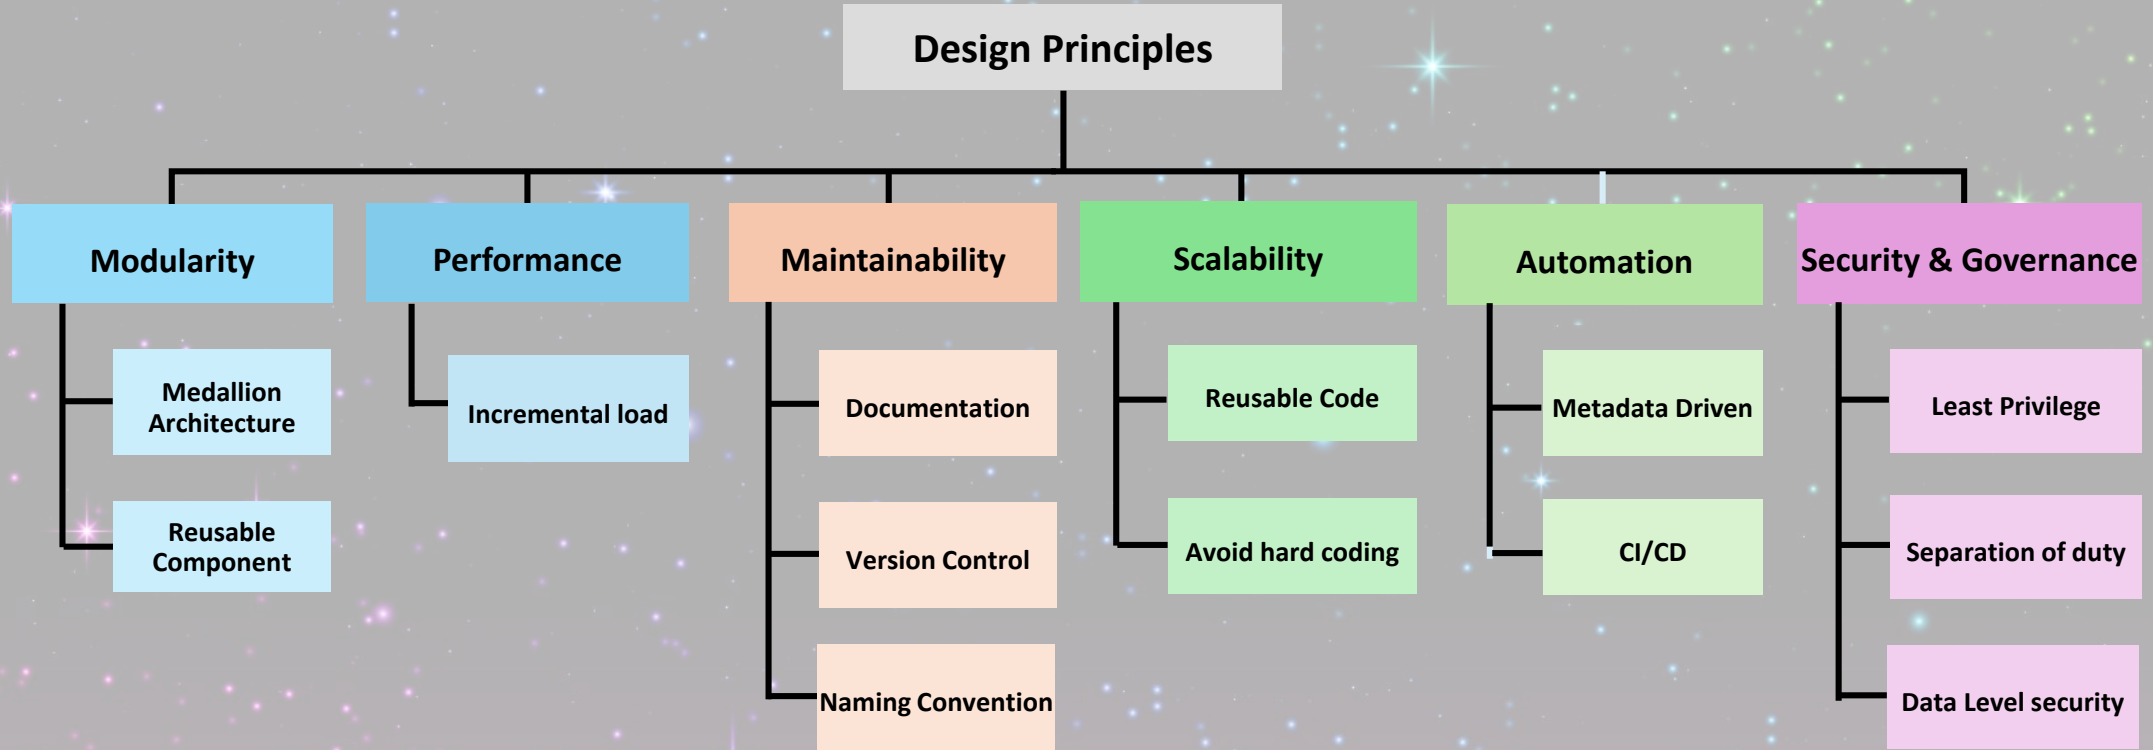
Fabric Terminology

Like **OneDrive**, OneLake is designed to be the single place for all analytics data.



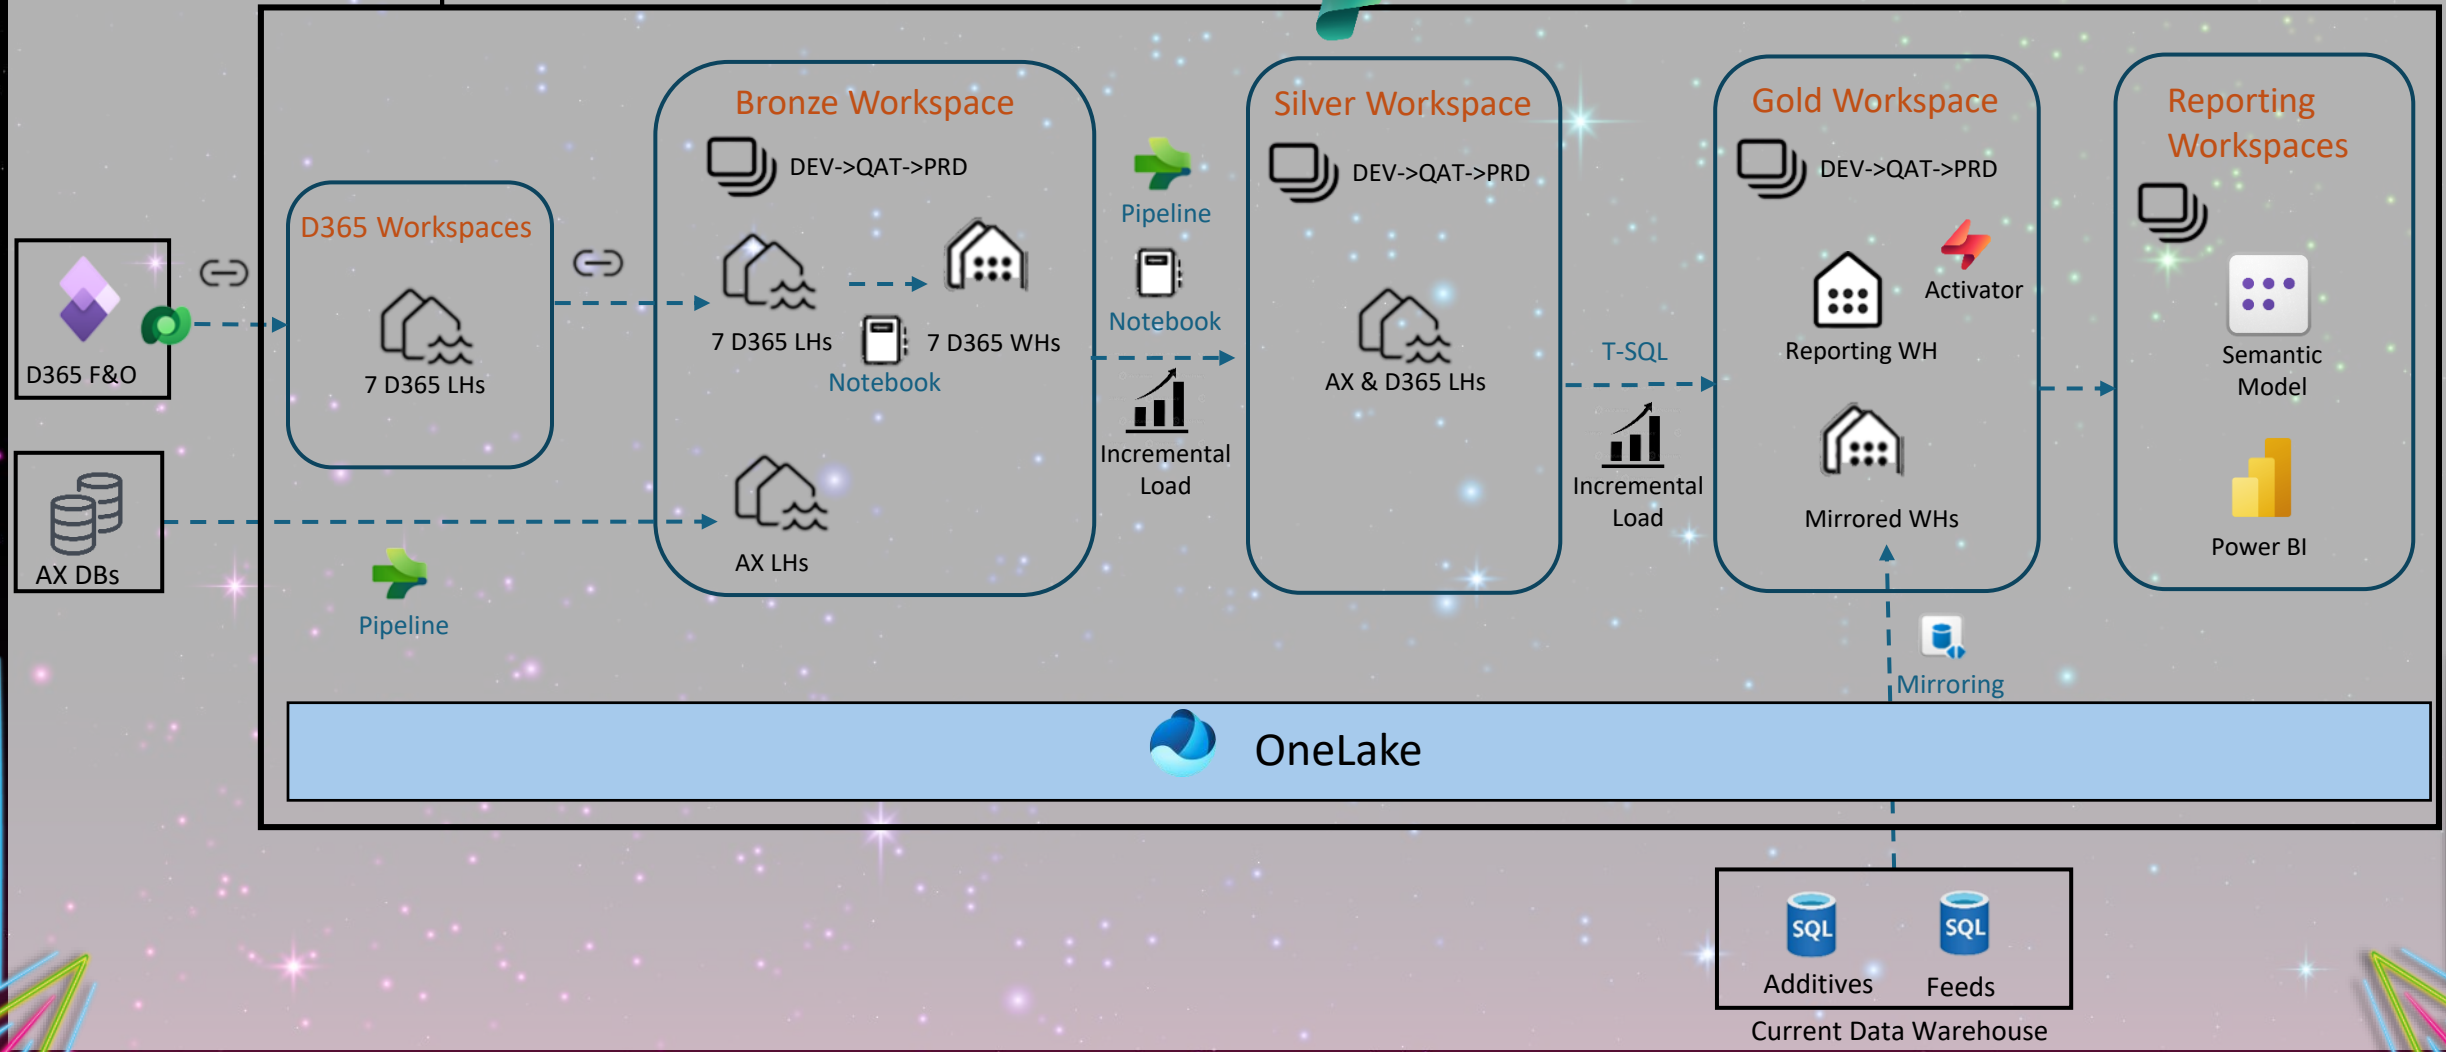
Core Design Principles

When we are building **Alltech Data Brain backbone Layer** for producing **clean, governed** data, we followed these core design Principles



Medallion Architecture

We built an Engineering **Backbone Layer** that ensures data is reliable, scalable and ready for business consumption.





Medallion Architecture

Bronze Layer: Raw data Ingestion

For example, Invoice line raw table from a certain AX or D365

	INVOICEID	INVOICEDATE	LINENUM	INVENTTRANSID	ITEMID	QTY	LINEAMOUNTMST	DATAAREID	RECID
1	INV000001	2015-10-04 00:00:00.000	1.0000000000000000	LOT0009310	MISC SUPP	18.0000000000000000	18.0000000000000000	USA	5637149185
2	INV000002	2015-10-04 00:00:00.000	1.0000000000000000	LOT0008560	24.2847.067.022	4000.0000000000000000	22360.0000000000000000	USA	5637149186

Silver Layer: Cleaned, Standardized data

Invoice Line table has all records from all AX and D365 instances with selected columns

	INVOICEID	ITEMID	QTY	SalesAmountHome	DATAAREID
1	CI000001	09.1041.067.175	50.0000000000000000	292.5000000000000000	GBR
2	8821	07.0225.067.077	-500.0000000000000000	-10950.0000000000000000	MYS
3	CI-00000001	01.4248.067.087	1000.0000000000000000	20500.0000000000000000	MYS
4	CI000001	HAULAGE	1.0000000000000000	18.2000000000000000	GBR
5	0020192	CON-008125	1.0000000000000000	109552.1400000000000000	BRA
6	CI-00000002	01.4241.067.087	1250.0000000000000000	26250.0000000000000000	MYS
7	CI-00000003	07.2301.067.087	50.0000000000000000	1645.0000000000000000	MYS
8	CI-00000004	02.1143.067.022	125.0000000000000000	6412.5000000000000000	MYS
9	F01-101-001138-CRI	24.0624.067.022	50.0000000000000000	170876.1200000000000000	CRI
10	CI-00000004	01.2198.067.087	625.0000000000000000	13375.0000000000000000	MYS

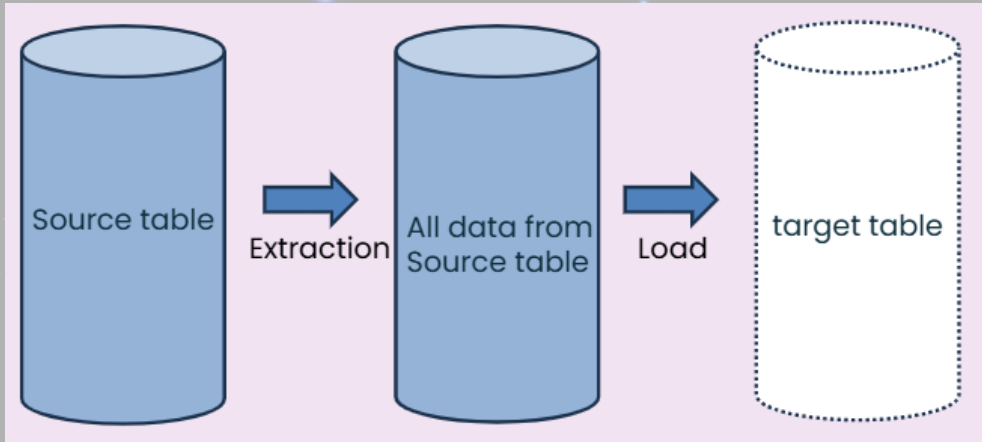
Medallion Architecture

Gold Layer: Integrated, Aggregated datasets optimized for reporting

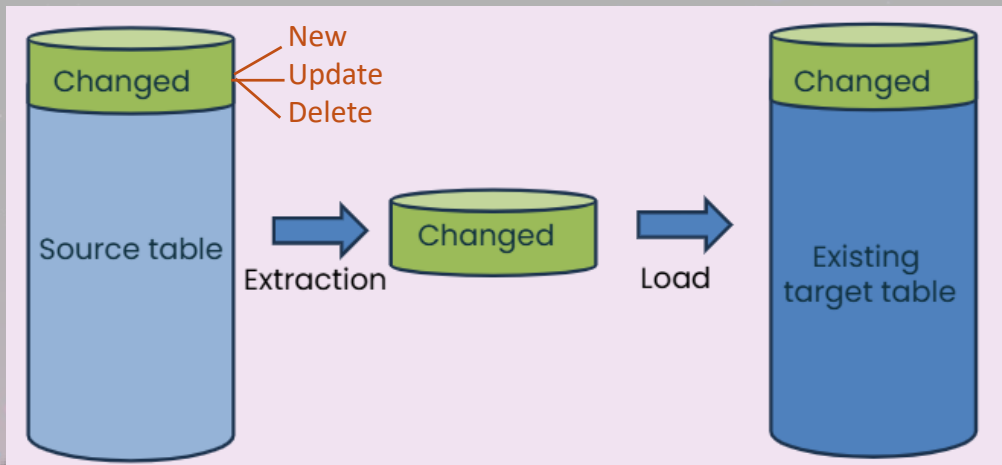


Incremental Processing

Full load:



Incremental load:



Aspect	Full load	Incremental load
Data Scope	Transfers the entire dataset	Transfers only new or modified data
Speed	Slow with large dataset	Fast
Data Volume	Small	Large
Resource Usage	Heavy consumption	Efficient
Complexity	Simple	Complex
Data Consistency	Guarantees absolute consistency	Requires logic for detecting and applying changes
Historical Data Restoration	harder	Preserves history

Metadata-Driven Framework

Metadata-Driven Framework

Bronze definition

- Load strategy
- Connections
- Table
- Variable Libraries
- Dynamic join conditions

Silver definition


- Connections
- Column mapping

Gold definition

- Dynamic view generation based on pre-defined join and derived conditions
- Curated data products (star schema + semantic models)

Power Apps for managing the Metadata Process

- Search or filter tables & columns
- Speed up process to build new data integrations
- Automated source metadata capture
- Include in Silver dropdown
- Table and column dropdown to reduce human mistakes
- Traceability (who made the change/addition, when and why)



What the framework provides

- Onboard new sources with metadata only.
- Auto type conversion to Fabric-ready data types.
- Standard load patterns.
- Unified auditing.
- Governance-ready: ownership, approvals, change history.

Metadata-Driven Framework

PowerApps created to manage metadata in Fabric



Selection_File

Search Table

Edit Data

Showing active records only

Id	Table Name	Column Name	instance_name	include_in_silver	ModifiedDate	ModifiedBy	Comment Re...	Dele
1	altglobalcoacons...	dataareaid		true	13/02/2026 12:29	Eashwar.dk@allt...	test	
2	altglobalcoacons...	globalconsolidat...		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
3	altglobalcoacons...	globalconsolidat...		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial loa	
4	altglobalcoacons...	id		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
10	altglobalcoacons...	IsDelete		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
5	altglobalcoacons...	recid		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
6	altglobalcoacons...	reversion		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
7	altglobalcoacons...	sinkcreatedon		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
8	altglobalcoacons...	sinkmodifiedon		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
9	altglobalcoacons...	sysrowversion		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
11	custinvoicejour	accountnum		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
12	custinvoicejour	cashdisc		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
43	custinvoicejour	CREATEDBY		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	
44	custinvoicejour	CREATEDDATETI...		true	04/02/2026 12:00	Artur.Ferreira@a...	Initial load	

← Home Enter New Record Details

D365 pdsbatchattribbyitem SinkModifiedOn

* Id table_name column_name

1262 pdsbatchattribbyitem SinkModifiedOn

instance_name include_in_silver * CommentsReason

true test adding new column

Save Record

← Manage the data here Search Table + Add New Record

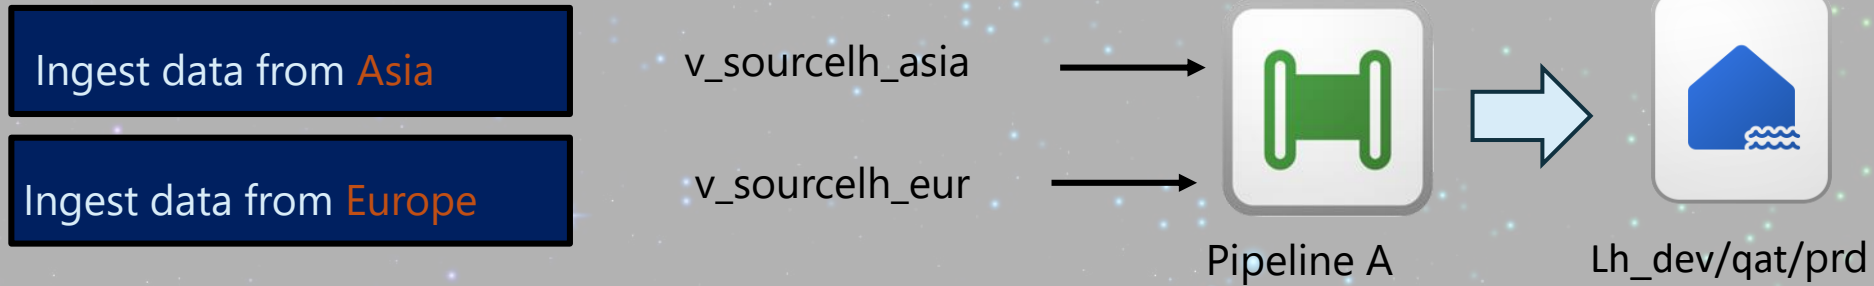
Showing active records only Selected ID: 1

Update	Id	Table Name	Column Name	Instance Name	Include in Silver
<input checked="" type="checkbox"/>	1	altglobalcoaconsolidation	dataareaid		true
<input type="checkbox"/>	2	altglobalcoaconsolidation	globalconsolidationacctname		true
<input type="checkbox"/>	3	altglobalcoaconsolidation	globalconsolidationcode		true
<input type="checkbox"/>	4	altglobalcoaconsolidation	id		true
<input type="checkbox"/>	5	altglobalcoaconsolidation	recid		true
<input type="checkbox"/>	6	altglobalcoaconsolidation	reversion		true
<input type="checkbox"/>	7	altglobalcoaconsolidation	sinkcreatedon		true
<input type="checkbox"/>	8	altglobalcoaconsolidation	sinkmodifiedon		true

Delete Save Changes

Modular & Scalable Design

Modular Pipelines aid reuse, leveraging the **Variable Library**



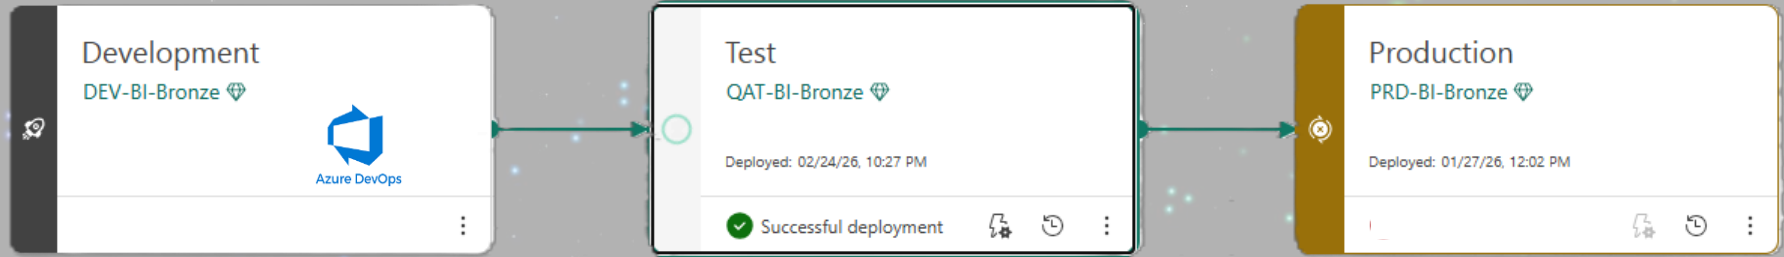
Variable Library

Var	DEV val	QAT val	PRD val
v_sourcelh_asia	asia_test	asia_prd	asia_prd
v_sourcelh_eur	eur_test	eur_prd	eur_prd
v_destdb	lh_dev	lh_qat	lh_prd

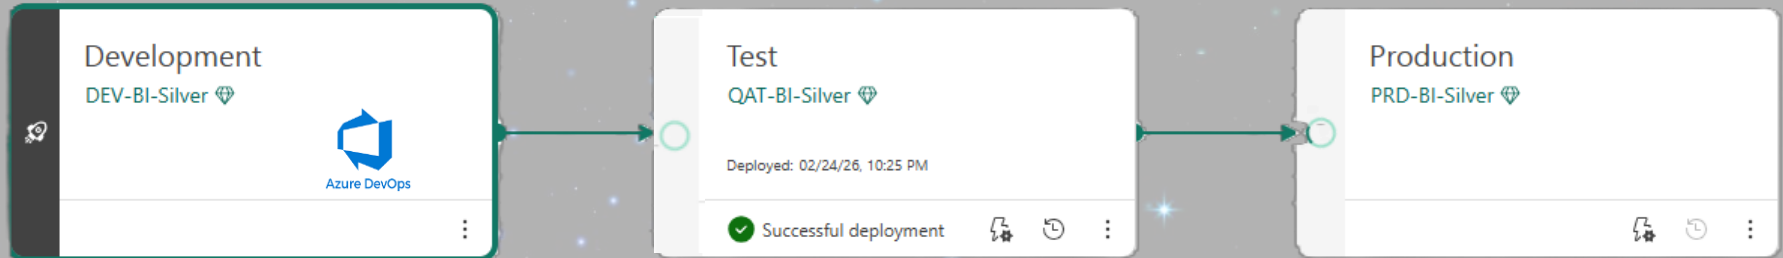
Fabric CI/CD



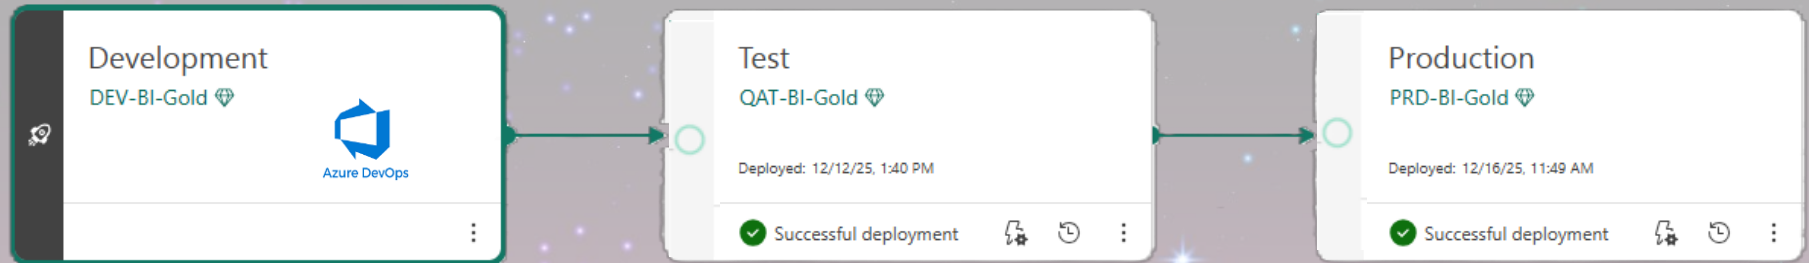
BI Bronze Pipeline



BI Silver Pipeline



BI Gold Pipeline



Domain Landscape

The 9-workspace structure is Data Engineering **backbone**, handled by IT/Data Engineering. It produces **clean, governed** data through Bronze, Silver, and Gold.

The **Domain Landscape** is your **business consumption layer**. It defines ownership boundaries for Finance, HR, Sales, and others.



The Gold Layer is the handoff point. Domain workspaces connect to it via OneLake Shortcuts. No Data is copied, domains simply read from Gold Layer.

Security & Governance



1 Finalize Security Architecture



Establish Layered Security Model

- Primary: OneLake security (data access)
- Secondary: Semantic model RLS (exceptions only)



Apply Column-Level Security

- Protect sensitive data (for example, cost)
- Reduce duplicate datasets by securing data upstream



Simplify Access via Personas

- Define standard personas
- Map personas to Active Directory groups



Define Data Ownership & Stewardship



Enable Controlled Self-Service BI

- Define Sensitivity labeling
- DLP enforcement (Control data export, sharing, and misuse)
- Data catalog & lineage



Prepare for Purview Implementation

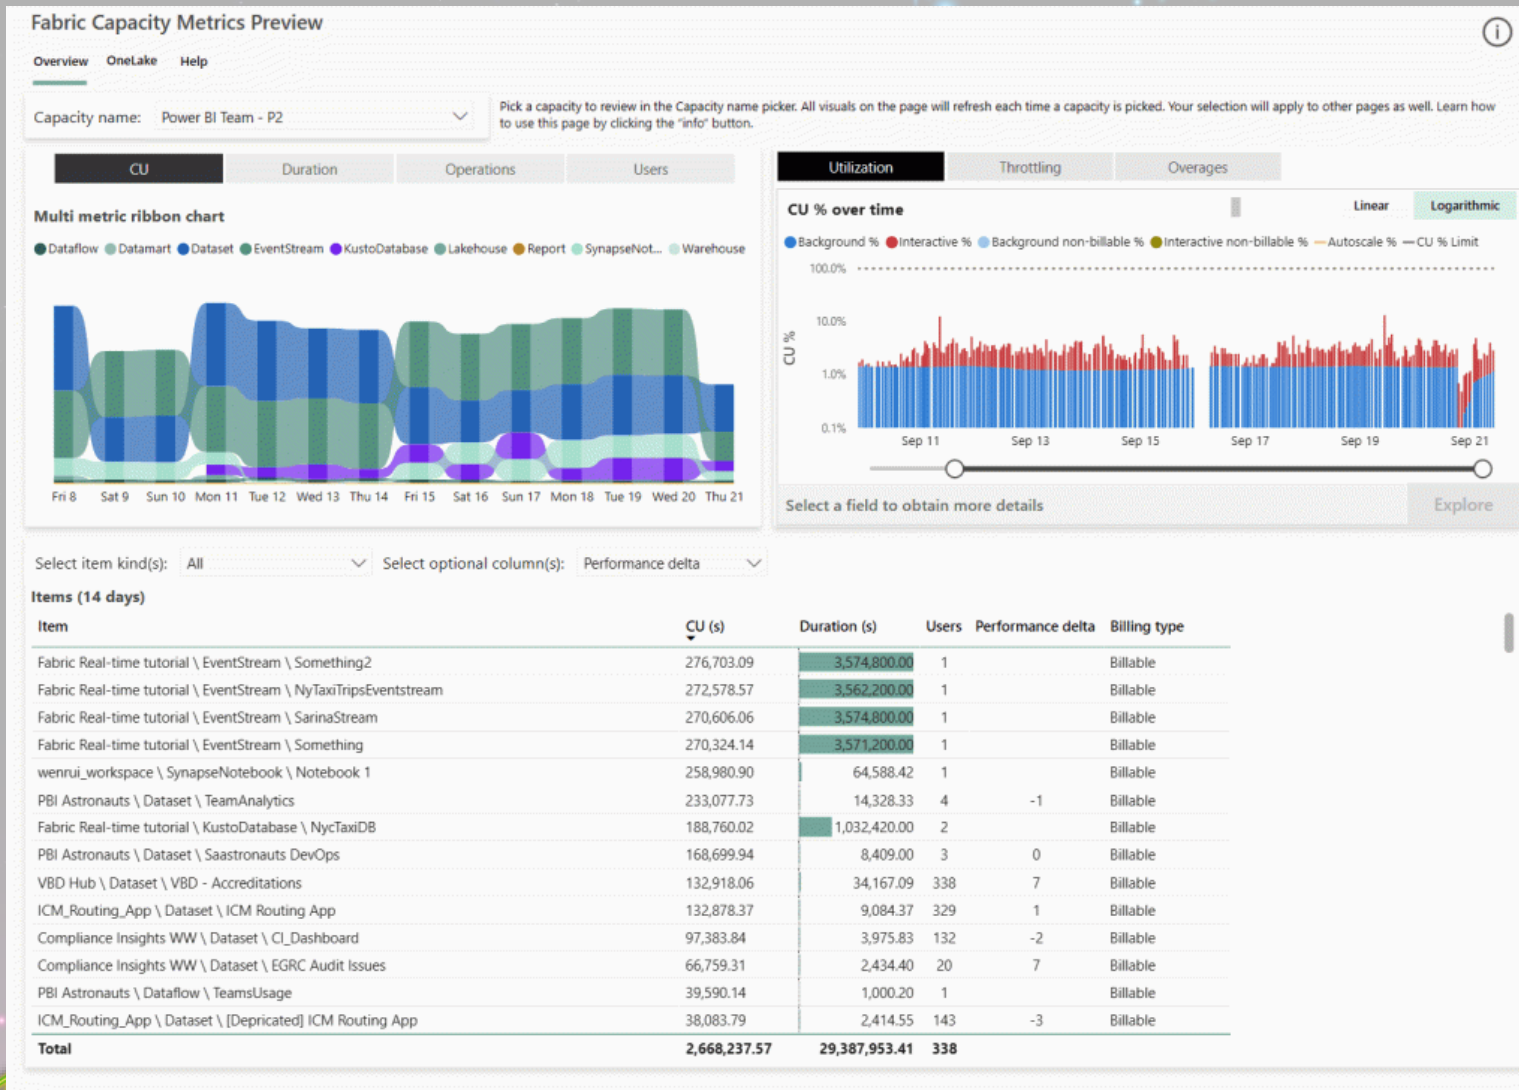
- Define Sensitivity labeling
- DLP enforcement (Control data export, sharing, and misuse)

Move the security setting **upstream** to the Lakehouse level

The screenshot shows the 'Reader' security settings page. At the top right is an 'Edit' button. Below the title, there are two tabs: 'Permissions' (selected) and 'Type'. Under 'Permissions', 'Read' is selected and 'Grant' is unselected. Below the tabs are two options: 'Data in role' (selected) and 'Members in role'. There is an 'Edit data' button and a search box. The main content is a table with columns 'Data', 'Type', and 'Data access'. The table lists a 'Schema' named 'dbo' and a 'Table' named 'publicholidays'.

Data	Type	Data access
Tables		
▼ [Icon] dbo	Schema	
[Icon] publicholidays	Table	

Performance & Cost Optimization



Performance
 Segmented by Capacity.
DEV&QAT: F32
PRD: F32

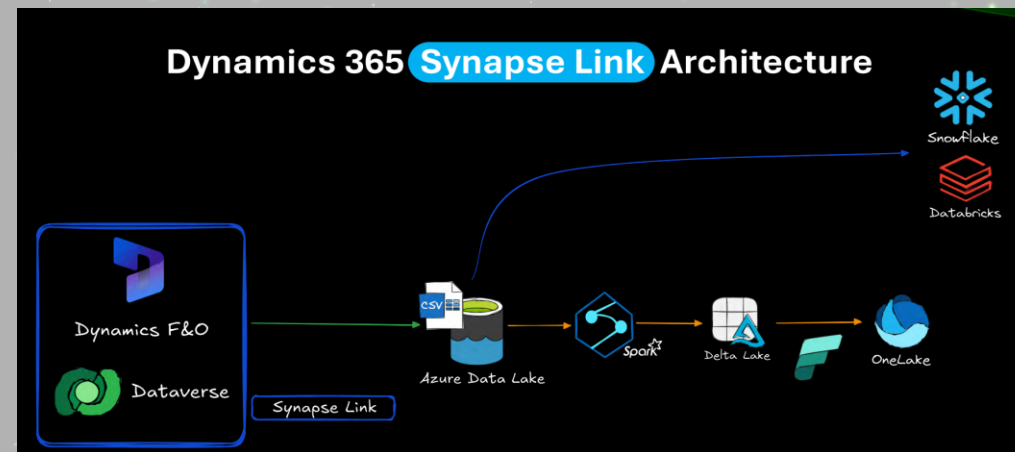
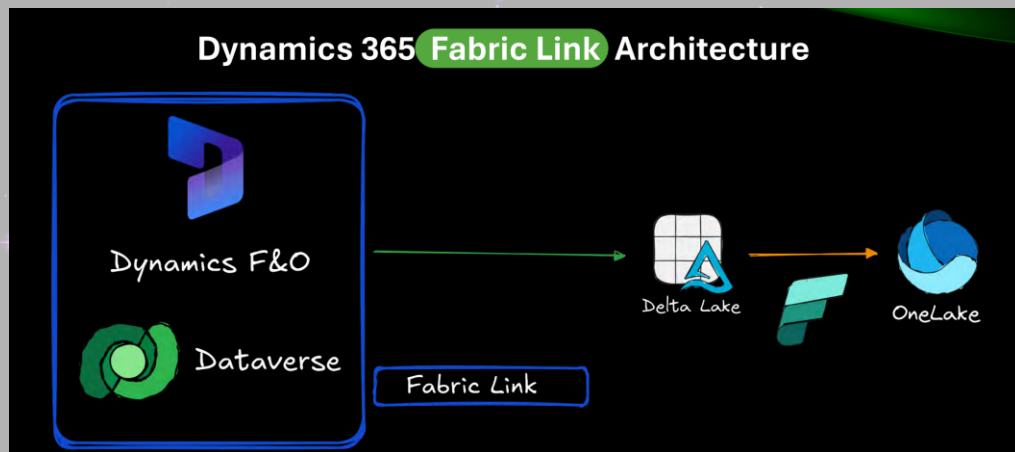
Cost Management

- Fabric Capacity Metrics App
- Bursting and Smoothing

Challenges with Fabric Link



- Fabric Link vs Synapse Link



- **Issue 1:** DimensionFocusBalance Table

=> **Solution:** Workaround – Instead using Synapse Link

- **Issue 2:** APAC data extraction Slowness after switched to Fast Fabric Link

=> **Solution:** Switch to Synapse Link, Unlink and Relink Fabric Link, switch back to Fabric Link

- **Issue 3:** Missing rows

=> **Solution:** Unlink and Relink, continuous monitoring of key tables

General Ledger Reporting in Fabric

General Ledger Revenue Breakdown




9 AX Instances




7 D365 Instances



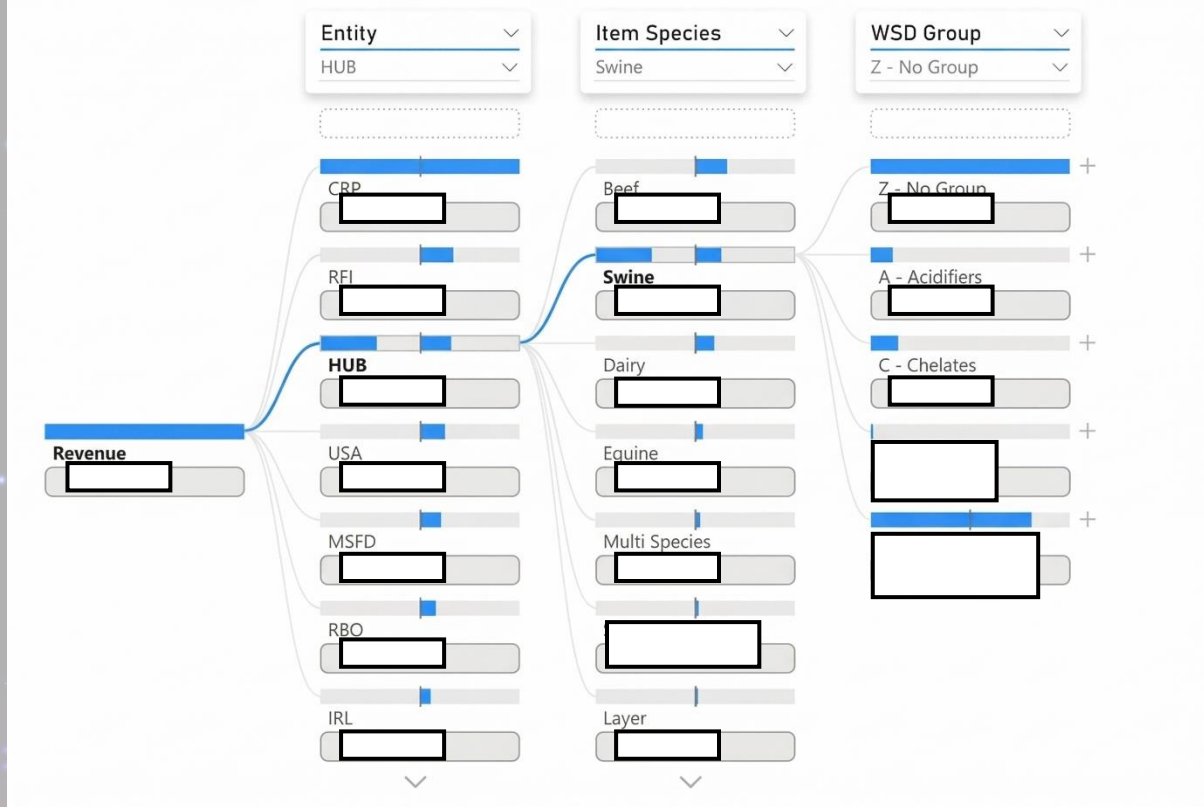
96 Entities



General Ledger data 2019 – Present Over 100M rows



General Ledger Revenue Breakdown D365



Process time : over 1 hour => 10 mins
Incremental load General Ledger data in GLD Layer
Full load twice caused throttling for the next 24 hours

Agenda



Power BI Intro

Power BI at Alltech

Power BI vs. Fabric

Power BI Service &
Reporting Structure

Self-service

Certification &
Sensitivity labels

AI & Data Agents

Putting it all together

Power BI Introduction



Power BI

- Connects to diverse data sources, transforms data with Power Query, models logic with DAX, and delivers interactive analytics

Power BI Suite

- **Power BI Desktop**
 - Windows app for report authoring and external tools
- **Power BI Service**
 - Cloud platform for sharing, collaboration, and online authoring
- **Other capabilities**
 - Mobile app, paginated reports, and on-premises reporting



Power BI Terminology

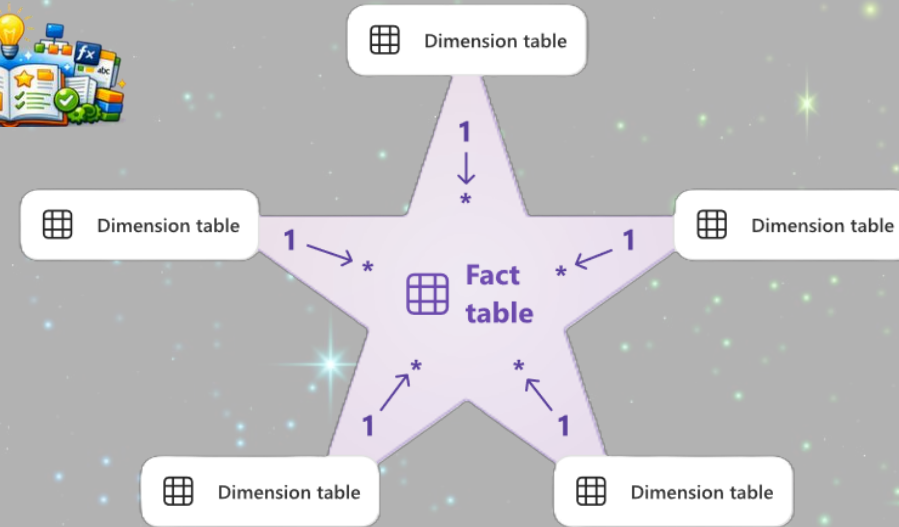


Semantic Model

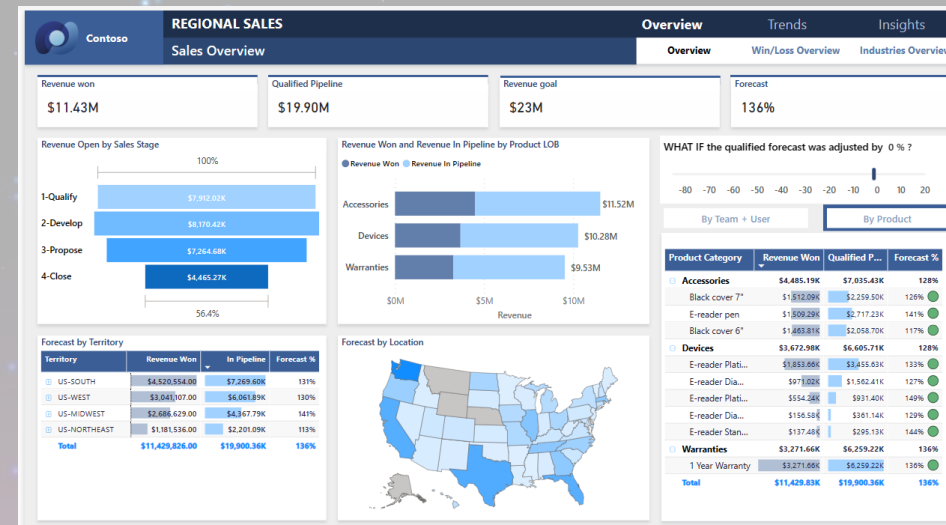
- Collection of related data tables for building reports and visuals (formerly called a dataset).

Report

- Visual insights from a semantic model
- Often called a "**dashboard**" colloquially, but a dashboard refers to a specific object in the Power BI



Example of star-schema semantic model



Sample report with visuals

Power BI at Alltech



The past

- **Excel-based** reporting caused inconsistencies due to varied data compilation
- **SSRS** and **Targit** improved accuracy but lacked strong permission controls
- Static reports were time-consuming to customize and update

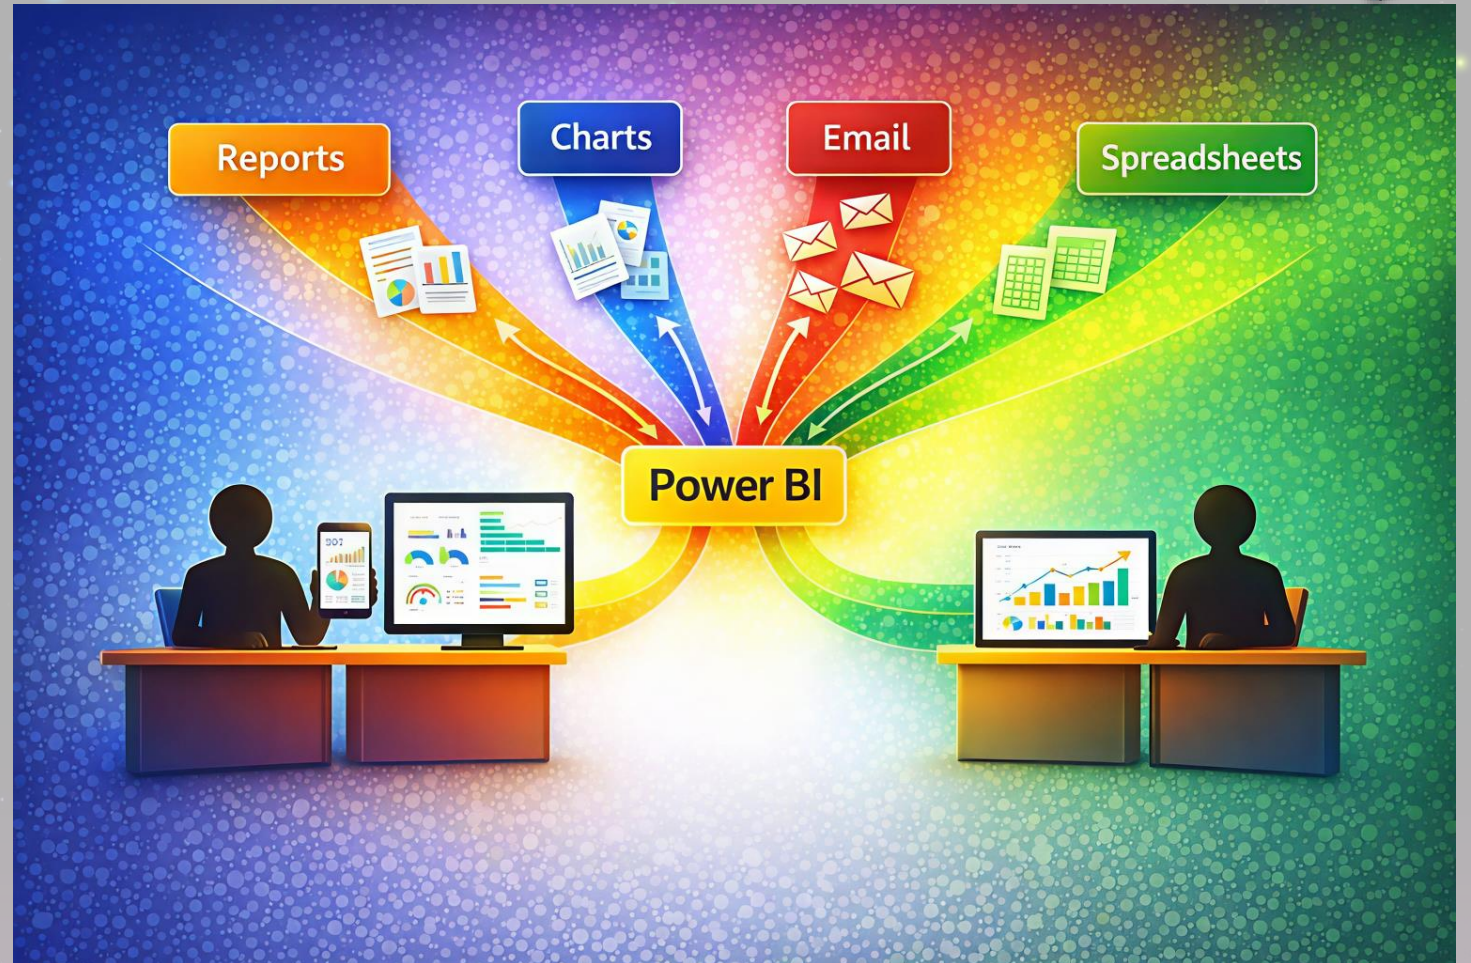


Power BI at Alltech



Switching to Power BI

- **Adopted Power BI in 2018** as part of Alltech's digital transformation to standardize data consumption
- Scaled from a small team to an **enterprise reporting platform** supporting automated insights
- **Prioritized reliable data foundations** with strong orchestration, security, and governance



Power BI vs. Fabric



- **Power BI** is now part of the **Fabric ecosystem**. While there are many overlapping features, there are some things unique to Power BI:



Feature	Description	Microsoft Fabric	Power BI
OneLake Integration	Unified data lake shared across workloads	✓	✓
Direct Lake Mode	Query OneLake data without import or caching	✓	✓
Notebooks & Dataflows Gen2	Advanced data preparation shared across Fabric	✓	✓
Data Activator	Real-time alerting and automation	✓	✓
Security & Governance	Centralized governance via Microsoft Purview	✓	✓
Copilot	AI-powered assistance for insights and report creation	✓	✓
Power BI Desktop	Windows app for report development	—	✓
Power BI Service	Cloud service for publishing, sharing, collaboration	—	✓
Paginated Reports	Pixel-perfect, printable reports	—	✓

Power BI – Home Cook Analogy



Imagine you are a **cooking dinner at home**:

- Buying groceries
- Preparing & seasoning
- Cooking
- Serving
- Eating

In a "**data**" sense:

- Getting data from sources
- Data staging & ETL
- Creating a semantic model & report
- Sharing semantic model and report to users
- Business users consuming reports



Fabric– Restaurant Analogy



Imagine you are operating a **restaurant**:

- Procuring ingredients
- Prep cooks preparing the ingredients
- Line cooks cooking the food
- Servers bringing food out to guests
- Guests eat the food, specifically what they **ordered**

In a "**data**" sense:

- Data Engineers getting data
- Pipelines and workflows to clean the data
- Power BI report developers making semantic models available
- Business users **consuming** the data that they want



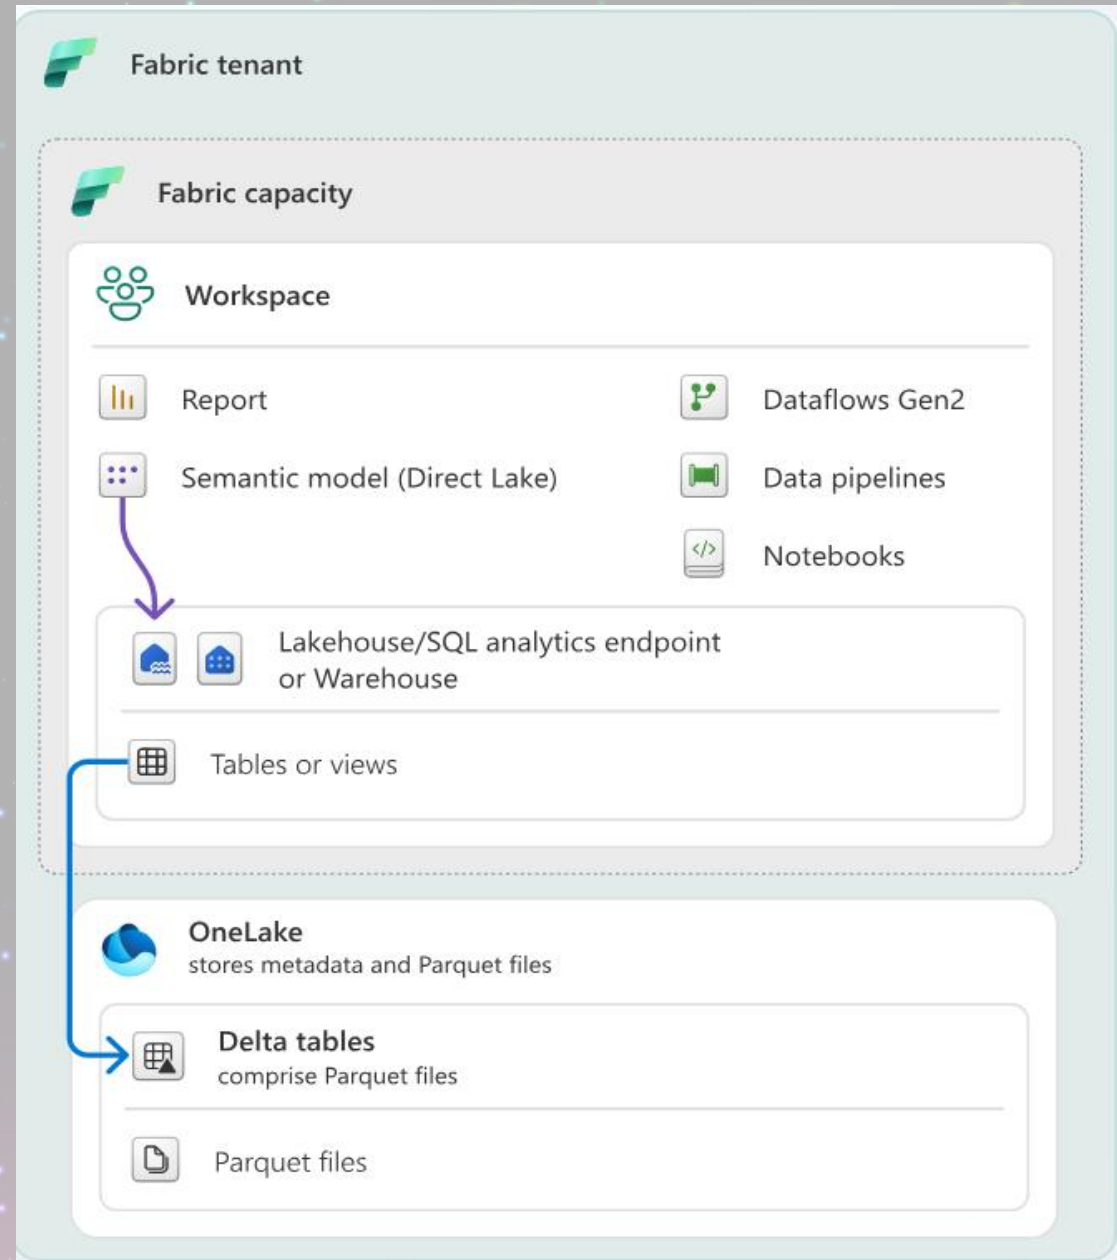
Direct Lake



Direct Lake enables Power BI semantic models to query data directly from **OneLake**

- Data remains in **Delta tables** with no duplication into the model
- Eliminates refresh requirements, simplifying operations and improving freshness

Framing gives model owners point-in-time control over what data is included in a semantic model



Reporting Structure - Workspaces



My Workspace

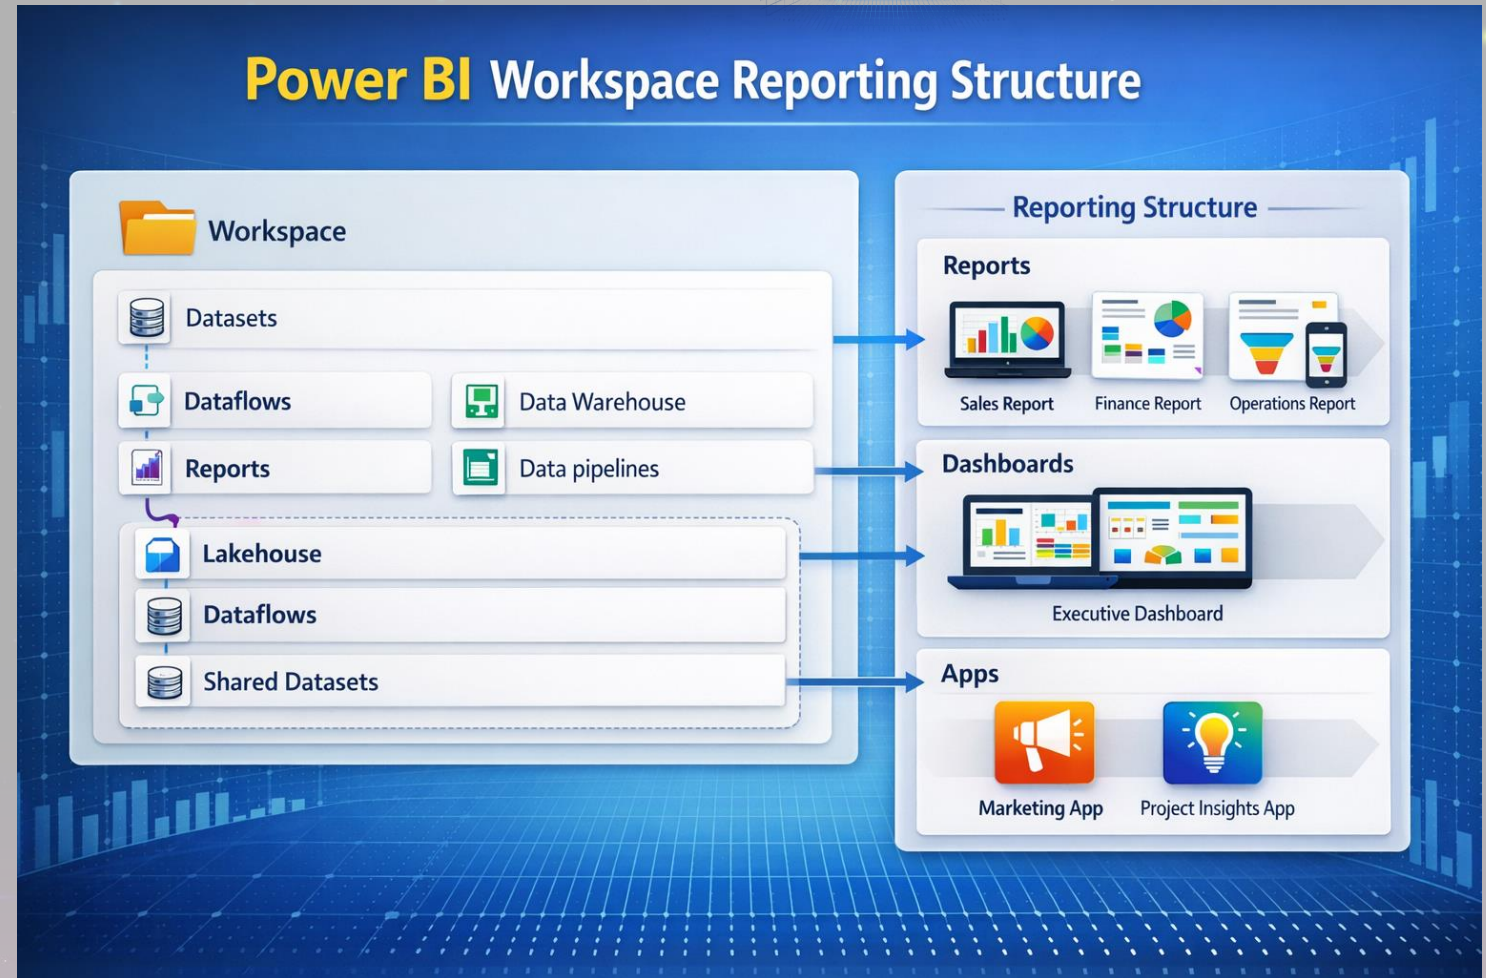
- Personal workspace for individual users
- Typically used only by report creators

Shared Workspaces

- Primary location for published reports and semantic models
- Support role-based access, Fabric capacity, and organized collaboration

Apps

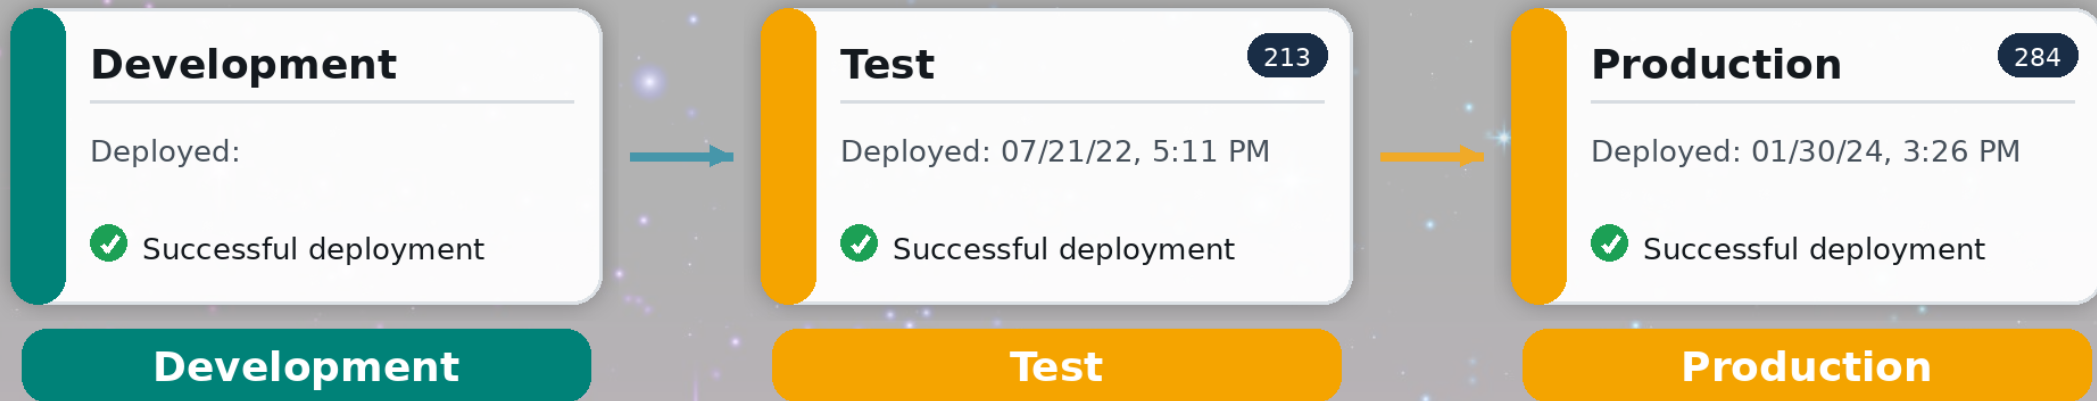
- Curated bundles of reports, dashboards, and models distributed to end users
- Can be automatically deployed and kept up to date



Reporting Structure – Workspace Strategy



- Workspaces can be organized by **department, region, business function, or project**
- A common pattern separates **distribution workspaces** (IT-managed) from **collaboration workspaces** (business-led with IT support).
- Distinct **Development, Test, and Production workspaces** support controlled delivery and governance
 - **Production workspaces** are the primary source for business consumption



Reporting Structure - Domains



- **Domains** define data ownership by grouping data around business departments
- Represent a **federated model**, shifting ownership from IT-centric to business-led
- **Workspaces** are assigned to domains, with all items included automatically
- Enable **domain-specific** settings to be applied consistently at scale
- **Business stakeholders own** their domains and related workspaces
- Support **subdomains** for finer-grained ownership and organization

IT

- Applications
- Infrastructure
- Service Management
- Security Logs

Finance

- GL / Accounting
- AP / AR
- Forecasting
- Budgeting

HR

- Workforce
- Payroll
- Learning
- Performance

Sales

- CRM
- Pipeline
- Orders
- Territories

Supply Chain

- Procurement
- Inventory
- Logistics
- Suppliers

Marketing


- Campaigns
- Customer Insights
- Digital Analytics
- Attribution

Power BI -Security



Power BI security is enforced through **access sharing** and **row-level security (RLS)**

- **Sharing permissions** control user access to semantic models and reports, including **Build permissions** for self-service analytics
- **RLS** restricts which data users can view based on defined roles and applies consistently to self-service use cases



Report



Business Users

Row-Level Security

Eastern US (0)

Members (0)

People or groups who belong to this role

Enter email addresses

Add

Fabric -Security



- **Fabric security** provides more granular control:
- **Object-level security (OLS)** restricts access at the individual item level, even without workspace roles
- **Column-level security (CLS)** limits access to sensitive columns, reducing the need for separate views
- **Row-level security (RLS)** controls which data users can see based on role definitions and supports **Direct Lake** semantic models.

All Lakehouse roles > Reader > Row and column security (Preview) ✕

Manage table level access for

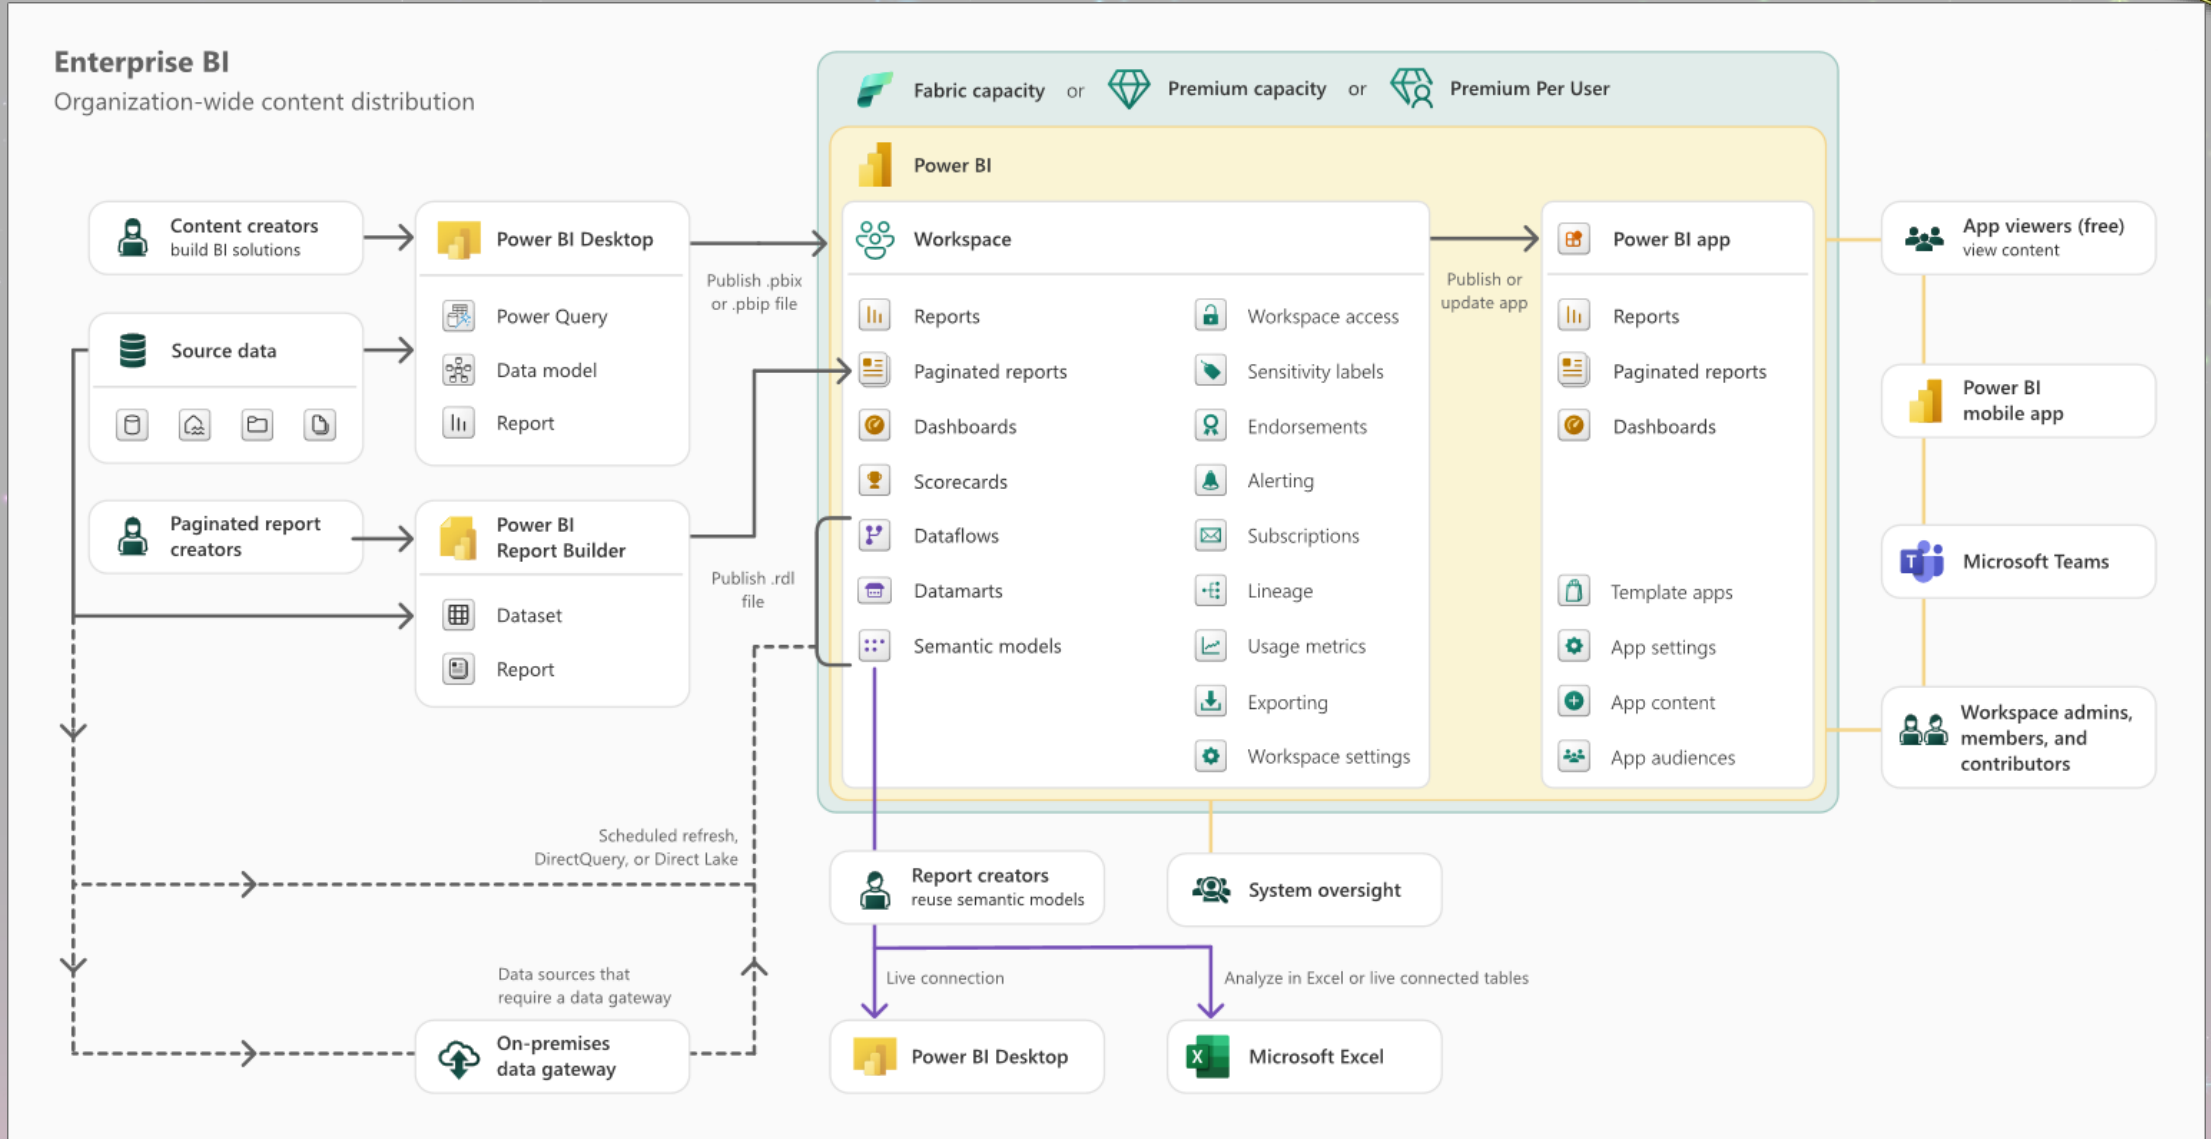
Manage access to your data by adding row and column level security to a table in a role. Members of this role will only be able to see specific data based on the rules you set below. [Learn more](#) 🔗

Row security (Preview) Column security (Preview)

Show data to members of your role if the following rules apply:

```
1 SELECT * FROM dbo.publicholidays WHERE
```

Power BI – Self Service



Power BI – Analyzing live data with Excel



Power BI enables Excel analysis

through live connections to governed semantic models

- Requires **Build permission**, with **row-level and object-level security** enforced
- Two primary options:
 - **Export to Excel (live connection)**: Explore a specific visual while staying connected to Power BI data
 - **Analyze in Excel**: Access the full semantic model in Excel for ad-hoc analysis using PivotTables

The screenshot shows the Excel interface in a browser window. The title bar reads 'Contoso Engineering Manufactur...'. The address bar shows a SharePoint URL. The ribbon includes 'File', 'Home', 'Insert', 'Draw', 'Page Layout', 'Formulas', 'Data', 'Review', 'View', 'PivotTable', and 'Share'. The main area displays a PivotTable with a data source labeled 'Engineering Manufactur...'. The 'PivotTable Fields' task pane is open on the right, showing a list of fields to be added to the report. The fields are organized into categories: 'better performance' and 'Fact_Production'. The 'Fact_Production' category includes fields like '%Availability...', '%OEE Rest', '%Performance...', '%Quality Rest', '%Target G...', 'Color Trans...', 'Current Rej...', 'dayNo1', 'Ideal Opera...', 'Lost Operati...', 'LostdayNo', 'Measure', 'Net Operati...', and 'NetdayNo'. The task pane also has sections for 'Filters', 'Rows', 'Columns', and 'Values'.

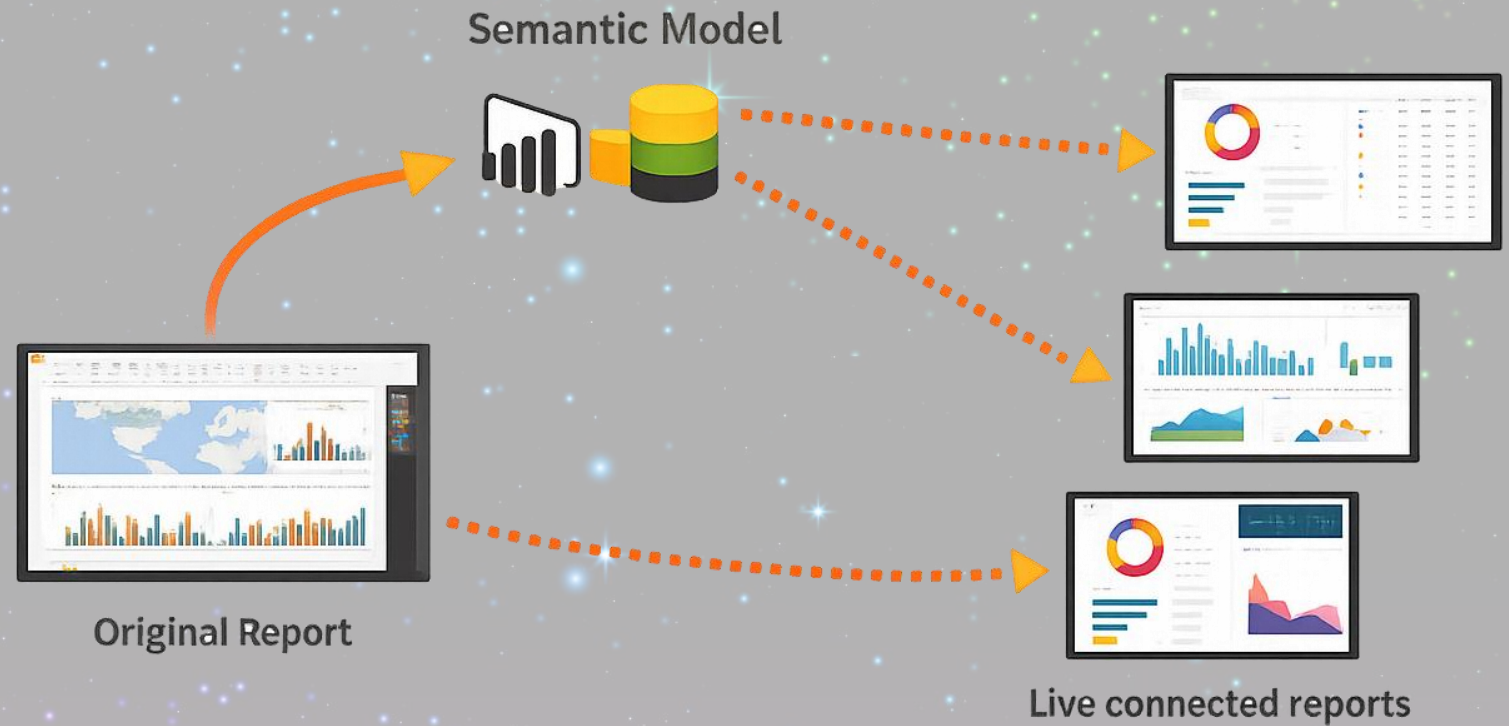
Example of Analyze in Excel PivotTable

Power BI – Connecting to Semantic Models



Power BI enables reports to **connect** to **published semantic models**, ensuring a single source of truth

- Reduces **duplicate and competing reports** across the organization
- Requires **Build permission**, with row-level and object-level security enforced
- Supports **live-connected report publishing** for secure sharing
- **Composite models** allow additional data sources to be extended on top of the governed model



Self-Service with Fabric



Visual query enables business user self-service for ad-hoc analysis. Business users don't need to worry that they accidentally change the underlying data



- Home
- Create
- Browse
- Workspaces
- Fabric_BI_DE
- TestLakehouse
- Data Engineering

Home

SQL endpoint

New SQL query New visual query

Explorer

+ Warehouses

- TestLakehouse
 - Schemas
 - dbo
 - Tables
 - test_concur_us...
 - Views
 - Functions
 - Stored Procedures
 - guest
 - INFORMATION_SCHEMA
 - sys

Data Query Model

Endorsement and Certification



- Power BI provides **two endorsement levels** to highlight trusted, high-quality data
- **Promotion:** Highlights valuable content and encourages reuse and collaboration
 - **Certification:** Indicates content meets organizational quality standards and is approved by **designated certifiers**

Improved discoverability:

- Makes it easier for users to find **endorsed semantic models**, even if they don't yet have access—reducing uncertainty around trusted data.

▲ Endorsement and discovery

Help coworkers find your quality content by endorsing this dataset and making it discoverable. [Learn more](#)

None
This dataset will appear in search results but isn't endorsed.

Promoted
When you're ready to distribute the dataset to your coworkers, promote it to let them know.

Certified
Certify your dataset to show coworkers that it's been reviewed and meets your org's certification criteria. [How do I get my dataset certified?](#)

Make discoverable
Allow users without access to this dataset to discover it and request permissions to access the data [Learn more](#)

i This dataset will be made discoverable. Others in your org will be able to find it by such details as name, tables, columns, etc. [Learn more](#)

Apply Discard

Sensitivity Labels



Sensitivity labels classify and protect organizational data without limiting collaboration

- Can be applied to **semantic models, reports, dashboards, and dataflows**
- **Labels are inherited** when data is accessed outside Power BI, including **Excel and PowerPoint**
- Can be applied in **Power BI Desktop, Power BI Service, and Fabric items**, with downstream assets automatically inheriting the label

The screenshot shows the Power BI web interface for a 'Finance Department' workspace. The 'Content' tab is active, displaying a table of items. The table columns are: Name, Type, Owner, Refreshed, Next refresh, Endorsement, Sensitivity, and Include in app. The 'Sensitivity' column shows labels like 'Confidential/Internal...' and 'Highly Confidential/In...'. The 'Include in app' column has a toggle switch set to 'Yes' for all items.

Name	Type	Owner	Refreshed	Next refresh	Endorsement	Sensitivity	Include in app
Finance Report	Report	Finance Department	10/29/19, 3:49:04 PM	—	—	Confidential/Internal...	<input checked="" type="checkbox"/> Yes
New Report	Report	Finance Department	6/5/19, 2:21:33 AM	—	—	Highly Confidential/In...	<input checked="" type="checkbox"/> Yes
Sales 2019	Dashboard	Finance Department	—	—	—	Confidential/Internal...	<input checked="" type="checkbox"/> Yes
Sales 2019 - External	Report	Finance Department	6/7/19, 6:32:05 PM	—	—	Confidential/Internal...	<input checked="" type="checkbox"/> Yes
T - Profitability	Workbook	Finance Department	6/5/19, 2:21:33 AM	—	—	—	<input checked="" type="checkbox"/> Yes

Example of sensitivity labels

Fabric Data Agents



Fabric Data Agents enable conversational Q&A over enterprise data using generative AI, similar to Copilot for Power BI

- Can be **instruction-driven**, ensuring responses align with specific business rules and use cases
- Support up to **five governed data sources**, including semantic models, lakehouses, warehouses, and KQL databases
- Can be **embedded across the Microsoft ecosystem**

The screenshot displays the Microsoft Fabric user interface. At the top, there's a navigation bar with 'Microsoft DXT', 'Fabric', and 'Gaming Trends Agent' tabs. Below this is a search bar and a 'Trial: 12 days left' indicator. The main content area is divided into two sections: 'Explorer' on the left and 'Test the agent's responses' on the right. The 'Explorer' section shows a tree view of data sources under 'Data', including 'GamingCustomerLH', 'Daily Support', 'Video game sales', and 'GameRatings'. The 'Test the agent's responses' section features a title 'Gaming Trends Agent' and a descriptive paragraph: 'This data agent acts as a virtual analyst for the gaming industry, providing data-driven insights to leadership and strategists. It can answer questions on user behavior, engagement, player profiles, sales, support issues, and game ratings. By integrating data from customer profiles, sales metrics, engagement analytics, support tickets, and sentiment sources like Reddit and YouTube, the agent produces detailed reports on games and market trends. Users can explore player demographics, performance across platforms and regions, customer satisfaction, and emerging support concerns to guide decisions. The agent is designed to help organizations understand their audience, optimize offerings, and identify opportunities in a fast-changing gaming market.' Below the text are three buttons: 'What are the historical trends across all my data?', 'Analyze recent data for any outliers', and 'Show me the details for a specific subcategory of data'. At the bottom, there's a text input field with the placeholder 'Ask a question to test the data agent's response' and a 'Sample questions' link.

Example of a Fabric Data Agent

Preparing Data for AI



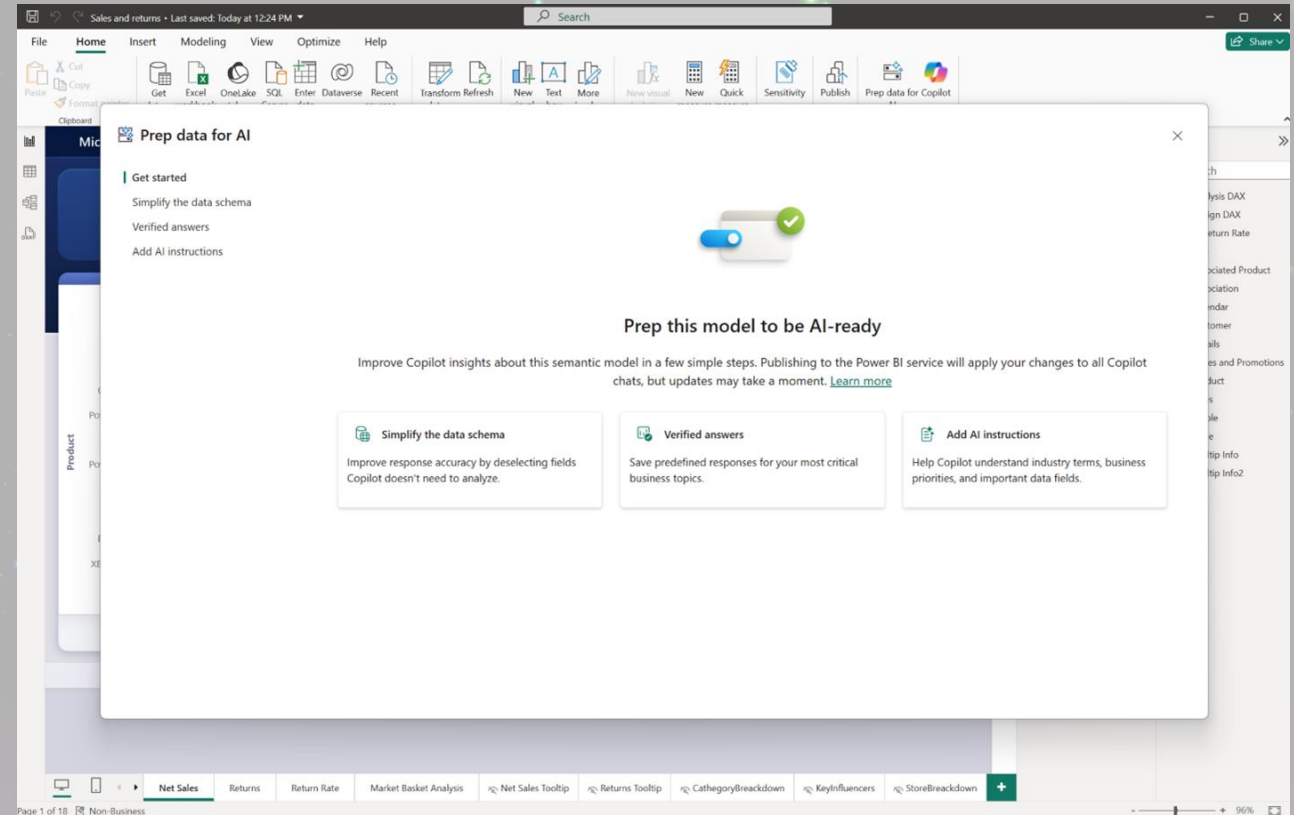
Preparing data is essential to maximize the value of **Copilot** and **Fabric Data Agents**

- Creates a foundation for **business-aware AI responses**
- Translates **business jargon into natural language**

Key preparation steps include:

- Limiting which fields are in scope
- Defining verified visuals
- Providing clear column and measure definitions

Once complete, data can be marked as **Prepped for AI**



Example of preparing data for AI

Ontologies



Ontologies define and unify business language across domains and OneLake data sources

- Represent **enterprise concepts** in a structured, consistent way

Define:

- **Entities** (e.g., Customer)
- **Properties** (e.g., customer name, email)
- **Relationships** (e.g., customers place orders through stores)

Once defined, ontologies are **bound to real data**, ensuring all downstream tools use the same language and meaning



Putting it all Together



Questions?



- **Thank you** for attending our session!
- We are happy to answer any questions!



Alltech's Fabric Journey: Building a Data Brain

Please fill out the survey and give us feedback!