Power BI developer life, reimagined with Fabric





Captain Platform

Thank you sponsors!

Platinum



Gold









Bronze











What to expect from this session

This is a feature-oriented overview session, not a deep technical dive/demo.

The goal is to provide a basic understanding of Microsoft Fabric and highlight how it impacts Power BI developers' workflows.





Anastasia SALARI

Principal Consultant at MDW, Switzerland

Passionate about Data

Microsoft MVP – Data Platform

Let's connect!

Survey





Microsoft Build 2023

(23.05.2023)

"Today we are unveiling Microsoft Fabric—an end-to-end, unified analytics platform that brings together all the data and analytics tools that organizations need. Fabric integrates technologies like Azure Data Factory, Azure Synapse Analytics, and Power BI into a single unified product, empowering data and business professionals alike to unlock the potential of their data and lay the foundation for the era of AI."

Captain Platform

“Is Fabric a Power BI ++?”

- Customer



Microsoft Fabric

The unified data platform for AI transformation

Data Factory

Analytics

Databases

Real-Time Intelligence

IQ

Power BI

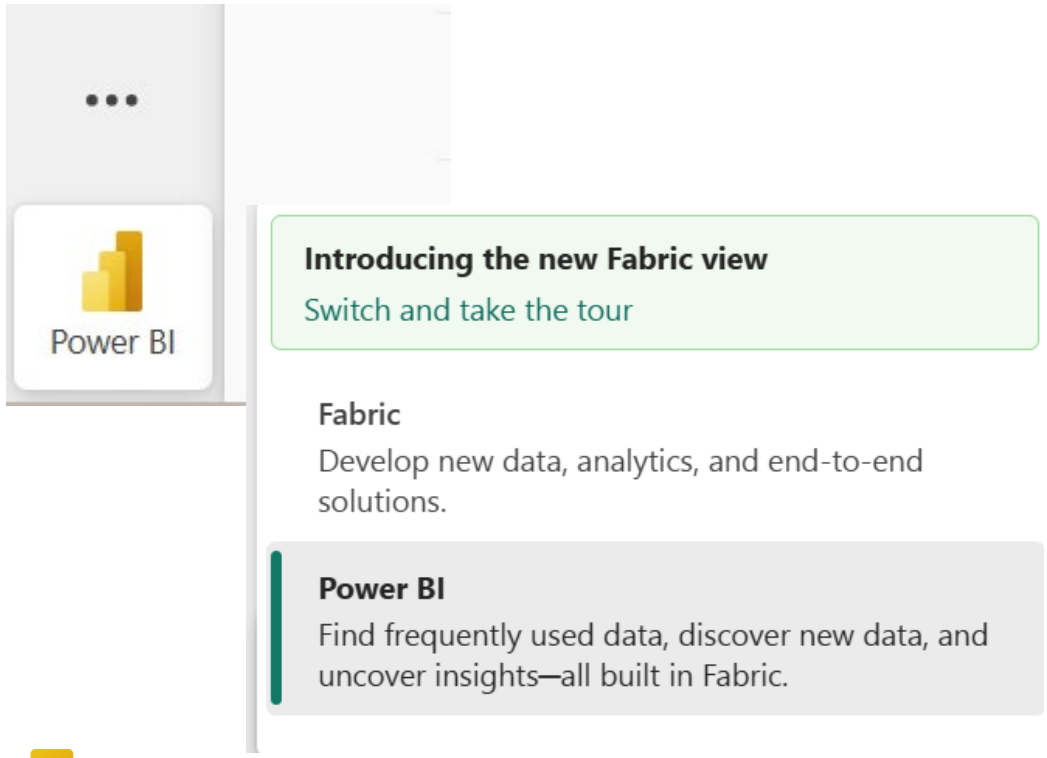
Fabric Platform

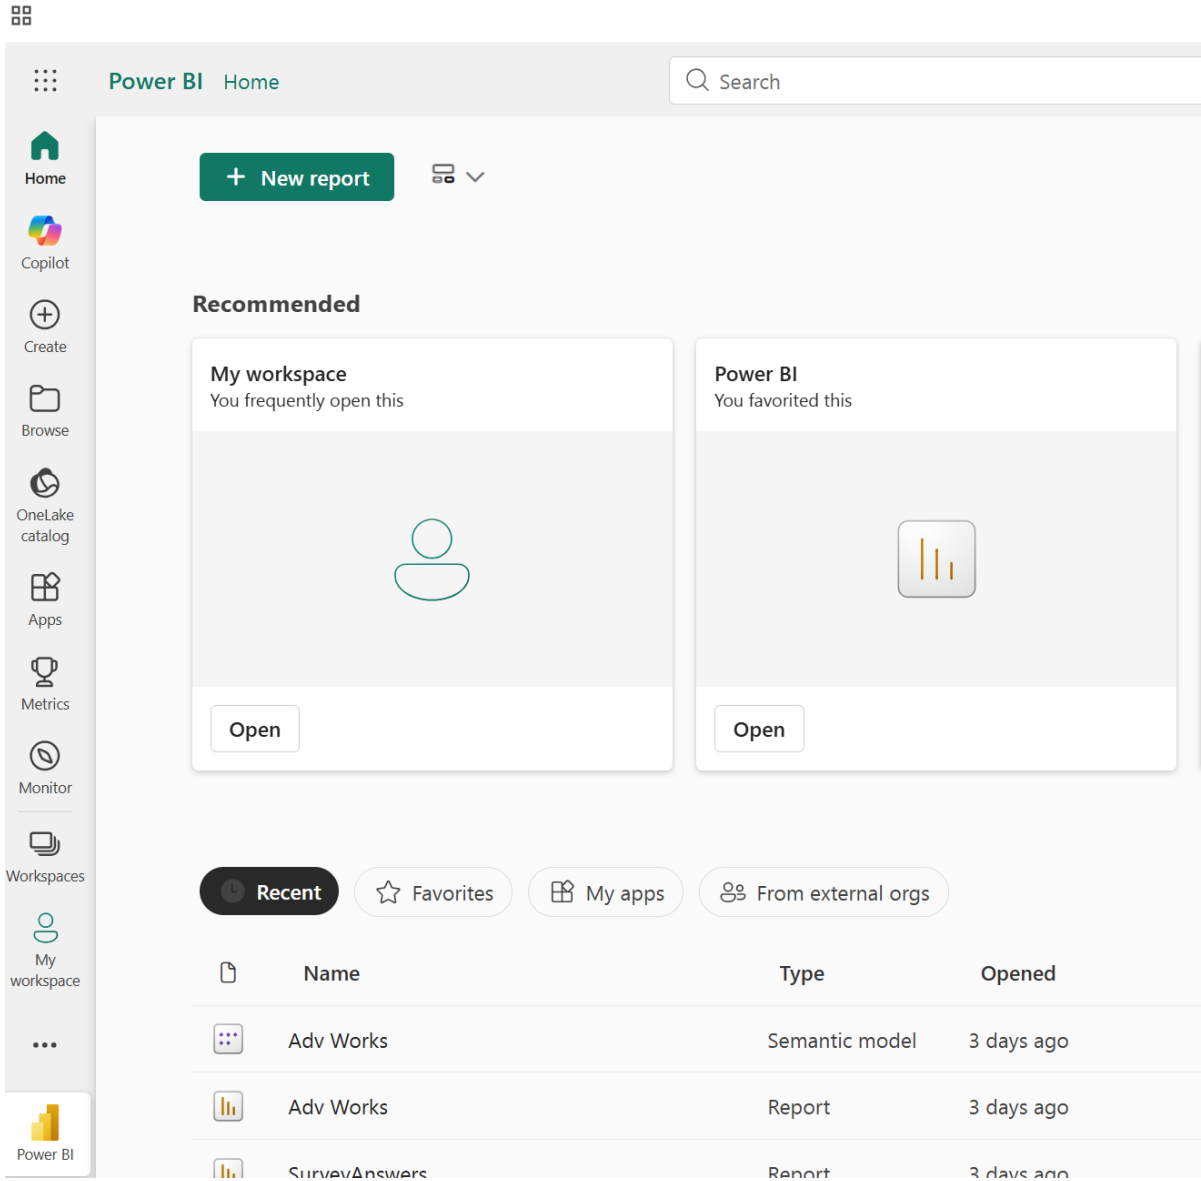
Copilot

OneLake

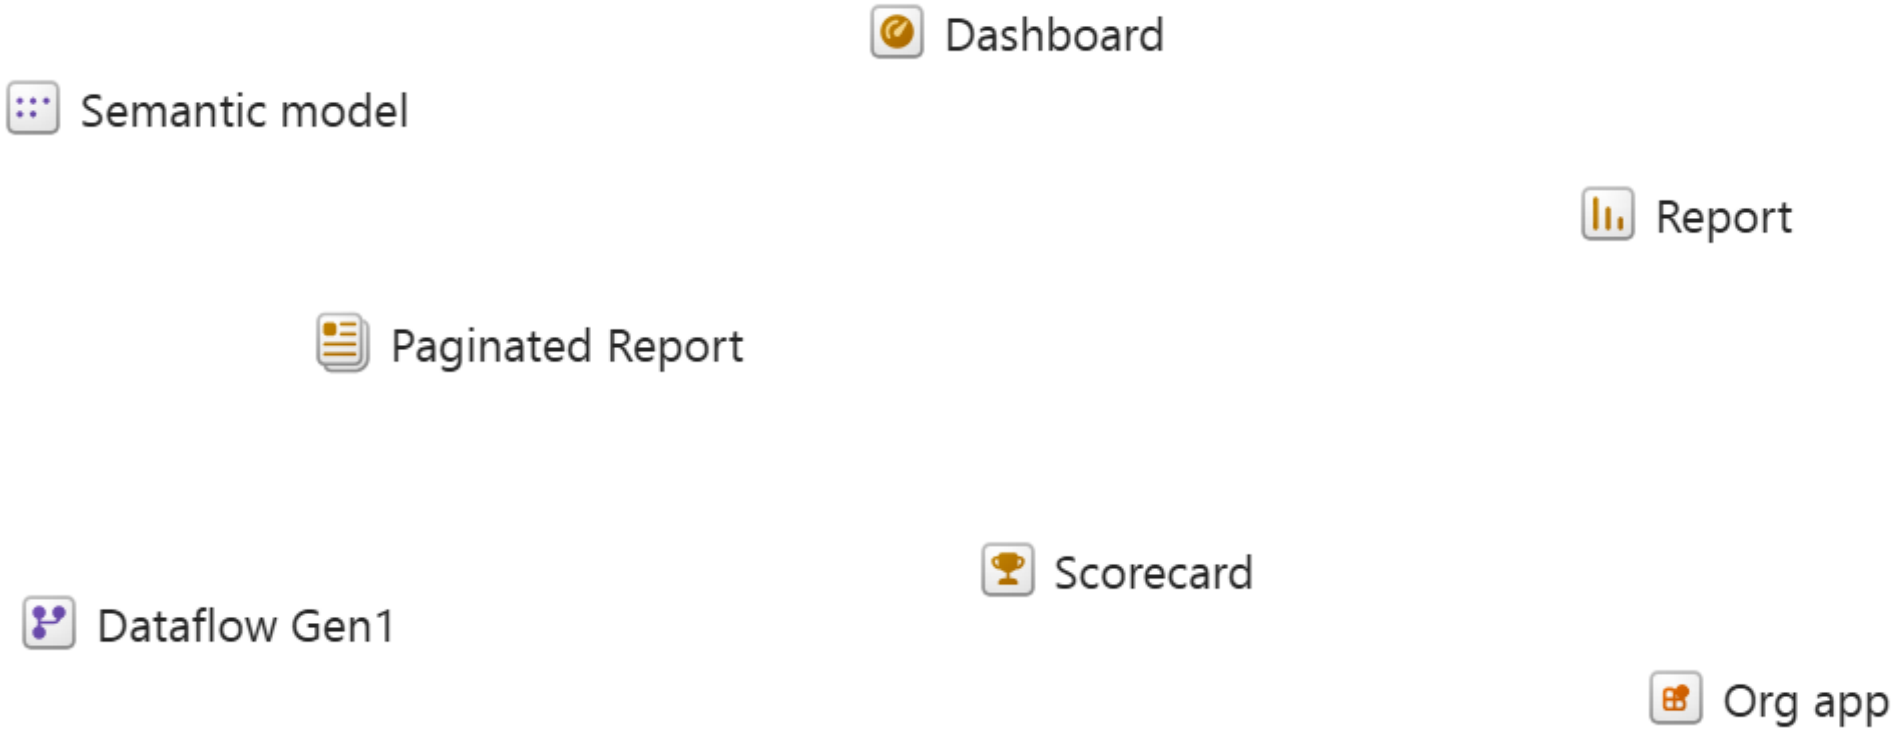
Governance

Power BI view



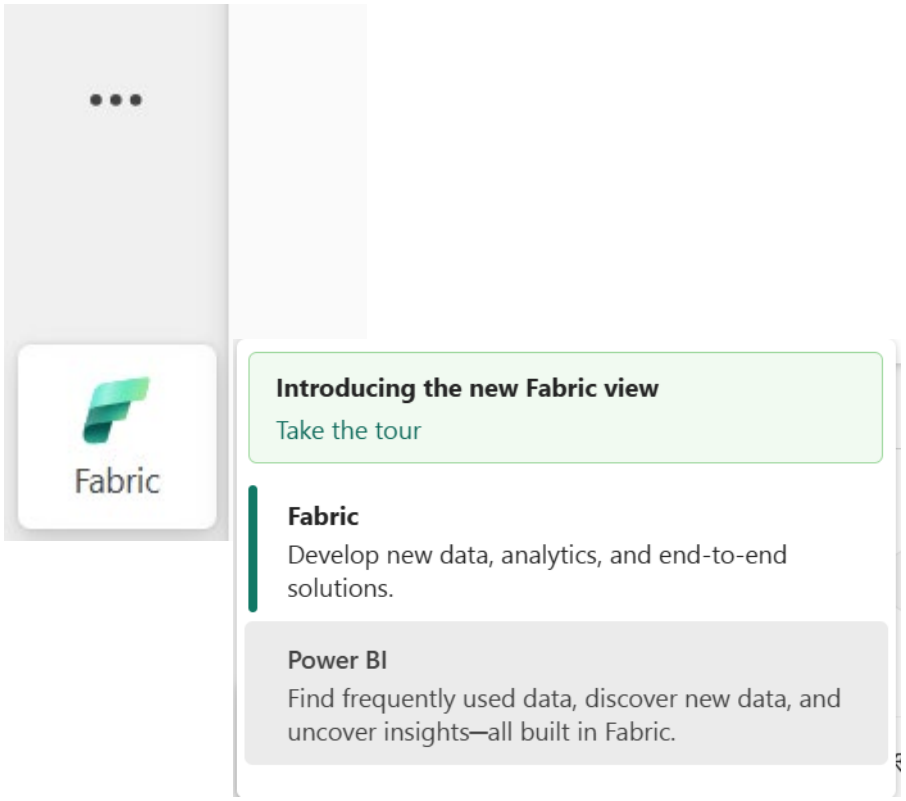


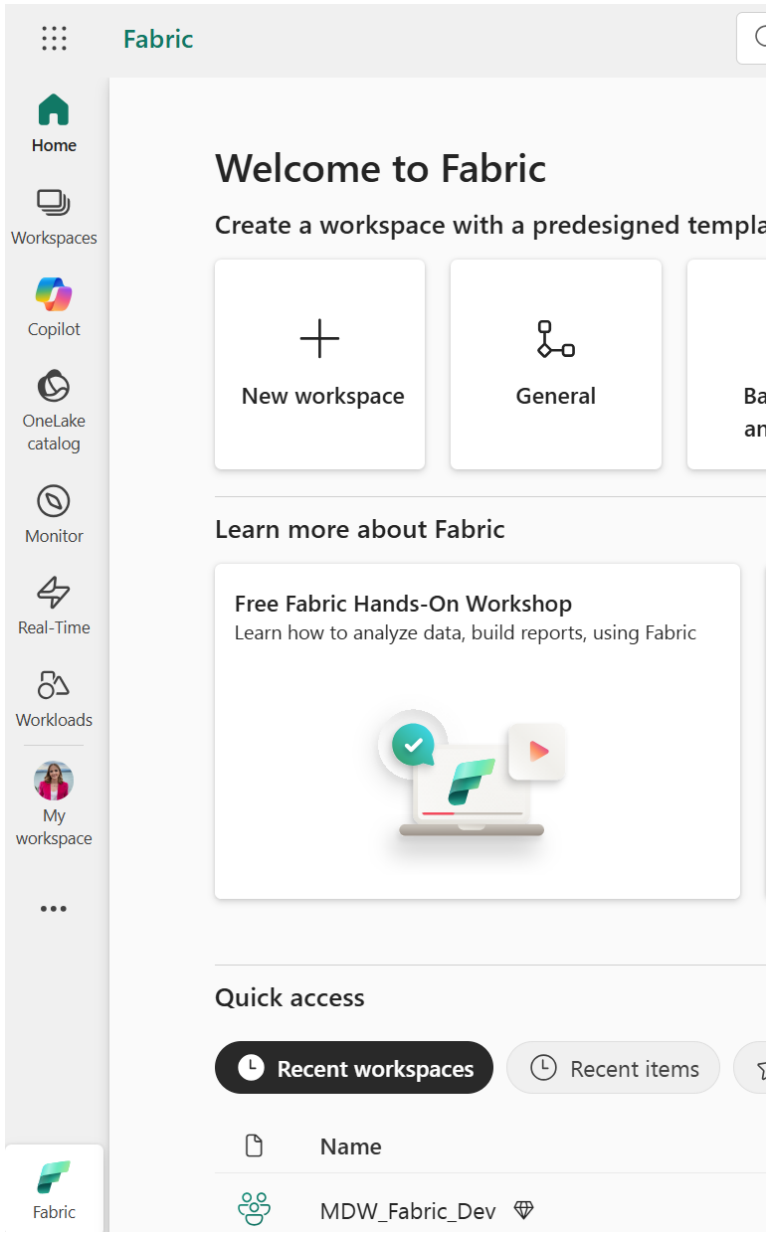
Power BI experience

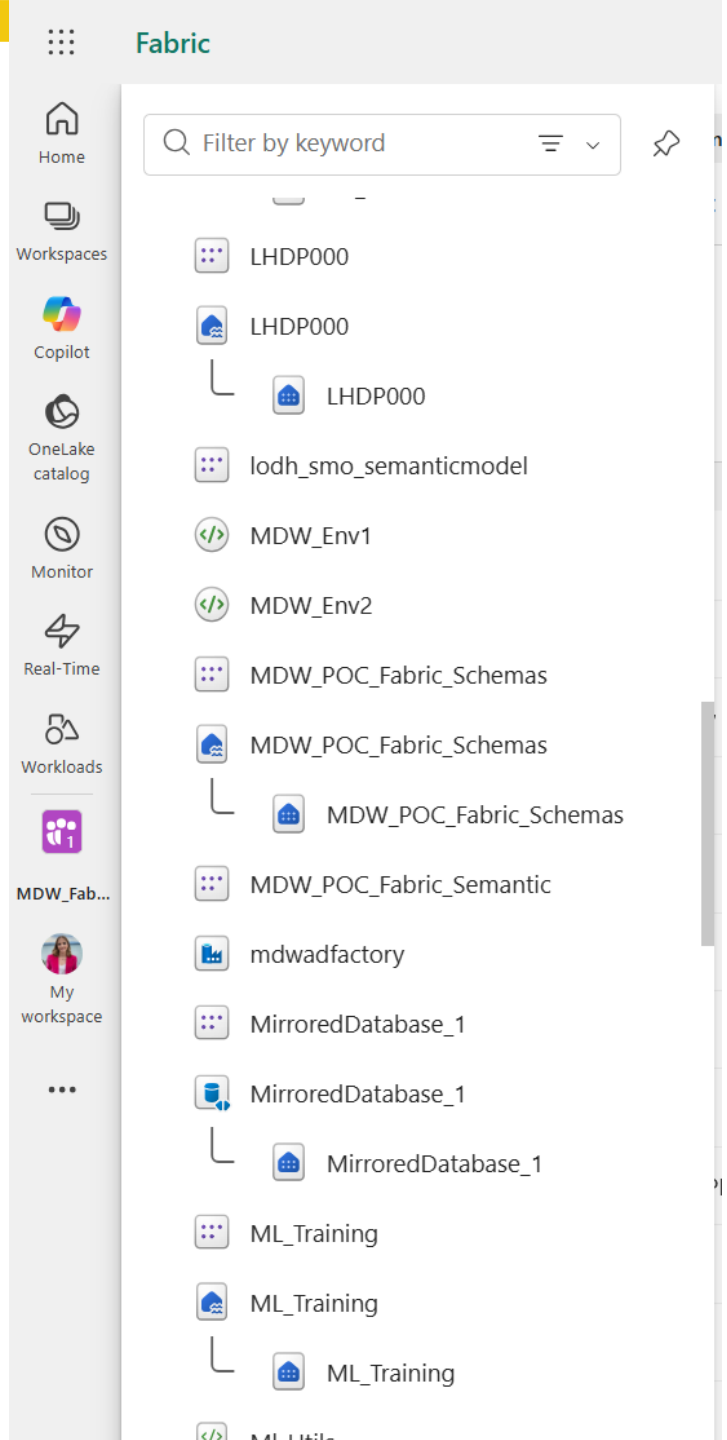


Captain Platform

Fabric view







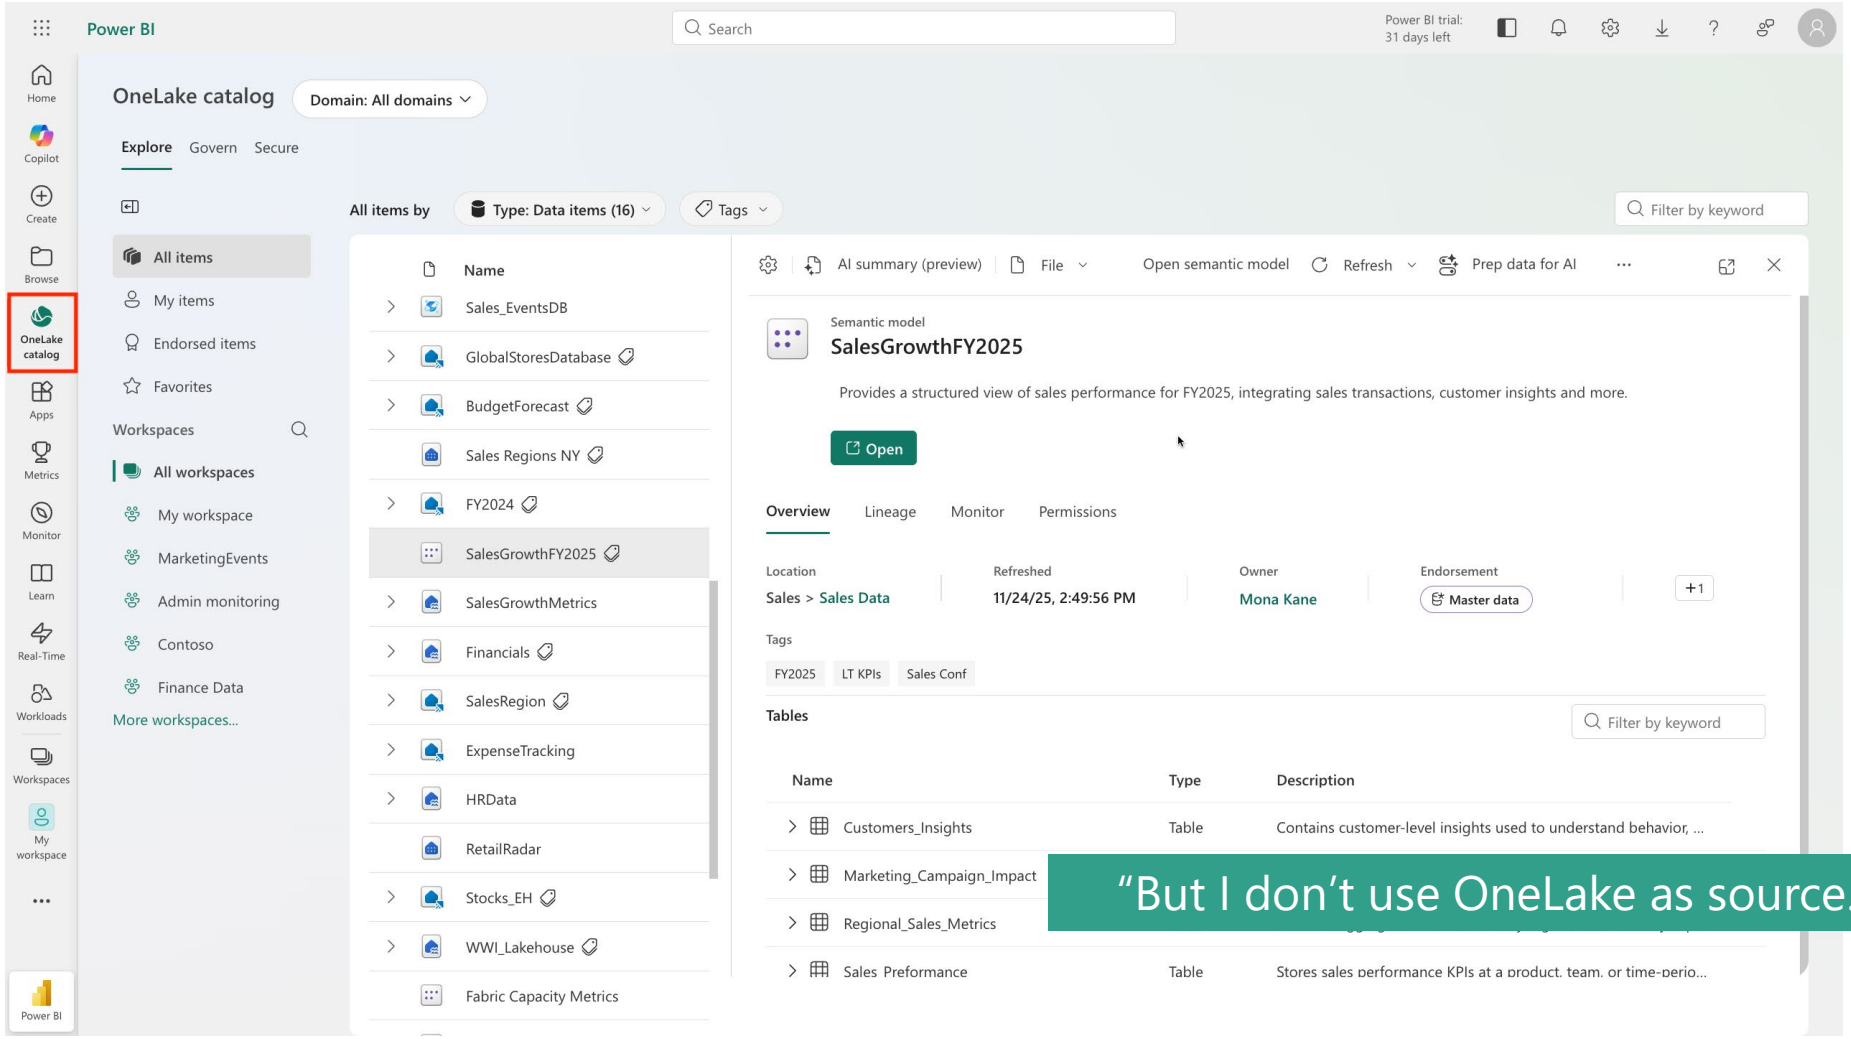
Captain Platform

Fabric

- Activator
- API for GraphQL
- Business Event Topic
- Custom stream connector
- Dataflow Gen1
- Datamart
- Digital Twin Builder Flow (preview)
- Eventstream
- Graph model
- KQL Queryset
- Mirrored Azure Databricks catalog
- ML model
- Operations agent
- Paginated Report
- Real-Time Dashboard
- Semantic model
- Spark Job Definition
- Variable library
- Anomaly detector
- Azure Data Factory
- Copy job
- Dashboard
- Dataflow Gen2
- dbt job
- Environment
- Experiment
- Graph queryset
- Lakehouse
- Mirrored database
- Notebook
- Org app
- Pipeline
- Report
- Semantic model (default)
- SQL database
- Warehouse
- Apache Airflow job
- Blueprint
- Cosmos DB database
- Data agent
- Dataflow Gen2 (CI/CD)
- Digital Twin Builder (preview)
- Eventhouse
- Exploration
- KQL Database
- Map
- Mirrored Dataverse
- Ontology
- Org app audience
- PostgreSQL Database
- Scorecard
- Snowflake database
- User data functions



OneLake catalog/hub



"But I don't use OneLake as source..." That's ok.

Building a dataset / report (classic way)

1. Power BI Desktop
2. Connecting to sources
3. (Optional) Chose among connection modes
4. (Optional) Do PowerQuery transformations
5. Defining tables, relationships, measures, etc.
6. Creating visuals, etc.
7. Publish to Power BI Service
8. Share with users

Store .pbix on local machine or in OneDrive/SharePoint



Classic workflow

Data/BI team using data tools

Data/BI team using BI/Reporting tools

How many of you could be in both groups/roles?



“Fabric” way

Data/BI team using data tools

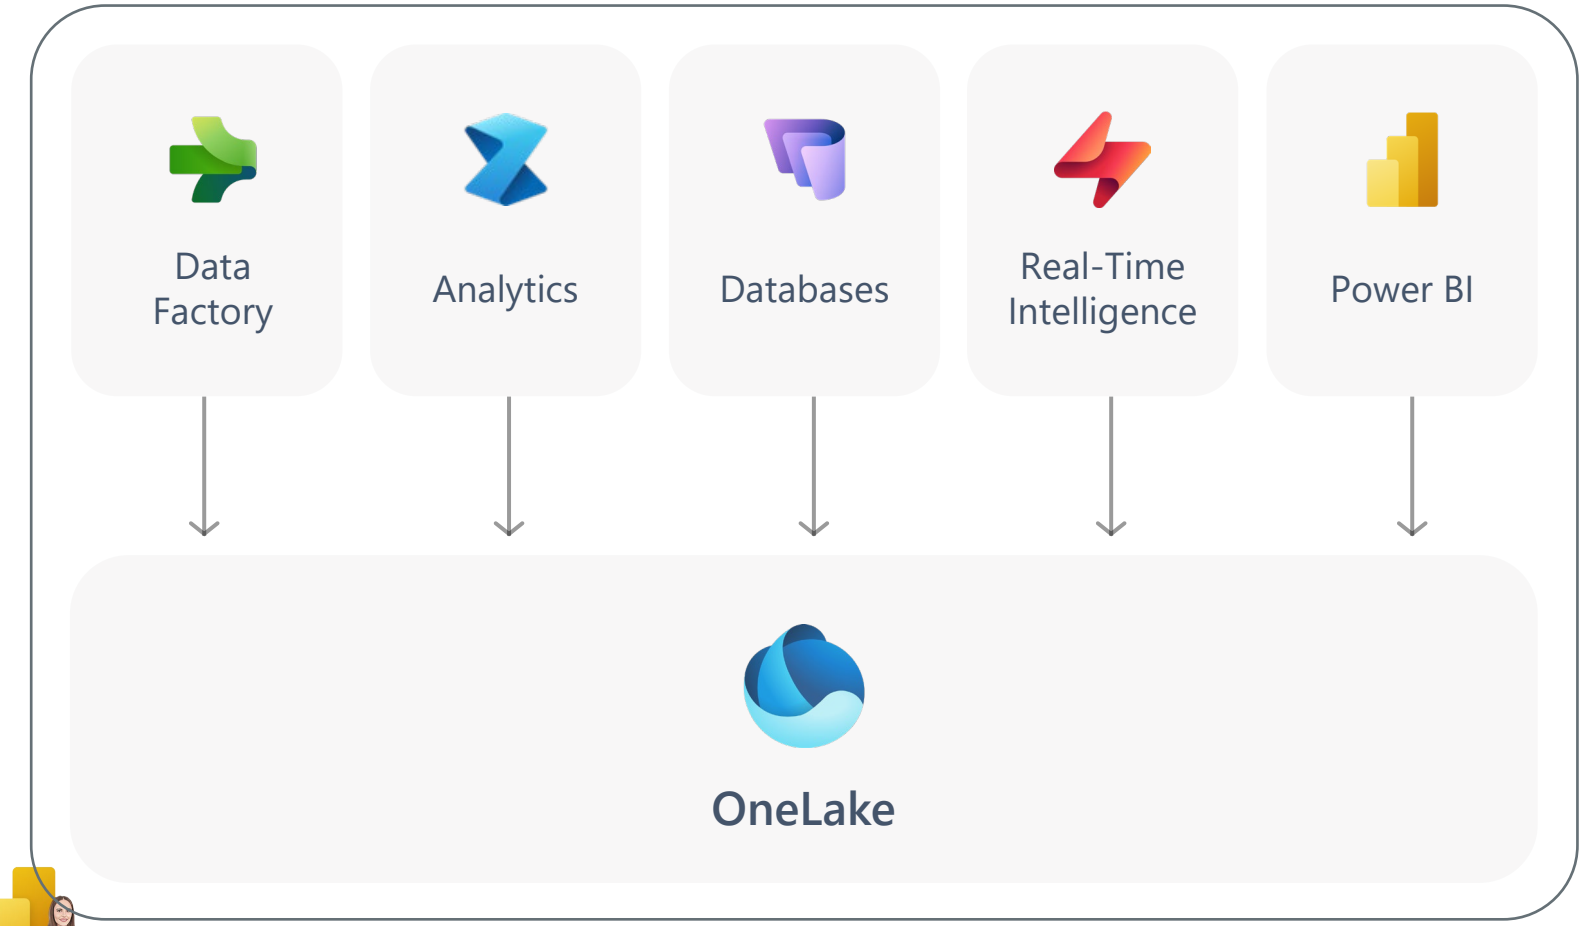
Data/BI team using BI/Reporting tools



OneLake

OneLake is the OneDrive for data

Your single, secure foundation for every workload



One open format for storing and querying data org-wide

Automatically indexed to simplify discovery and sharing

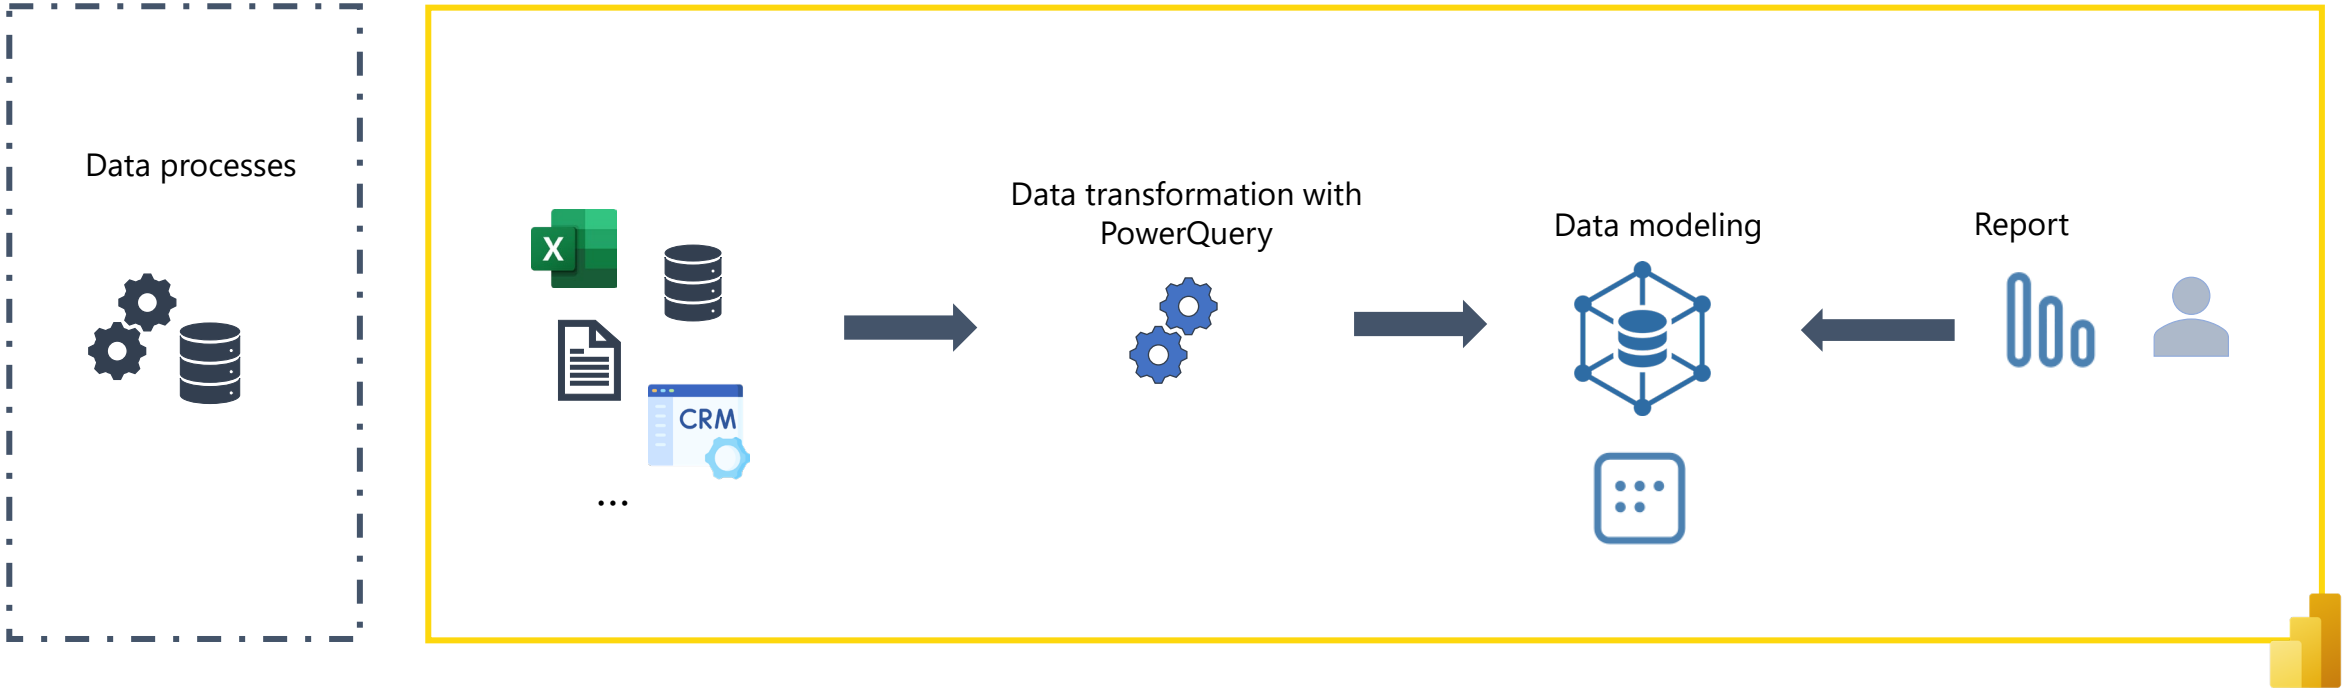
Built-in lineage for **end-to-end traceability** across workloads

Protected at every layer with OneLake security

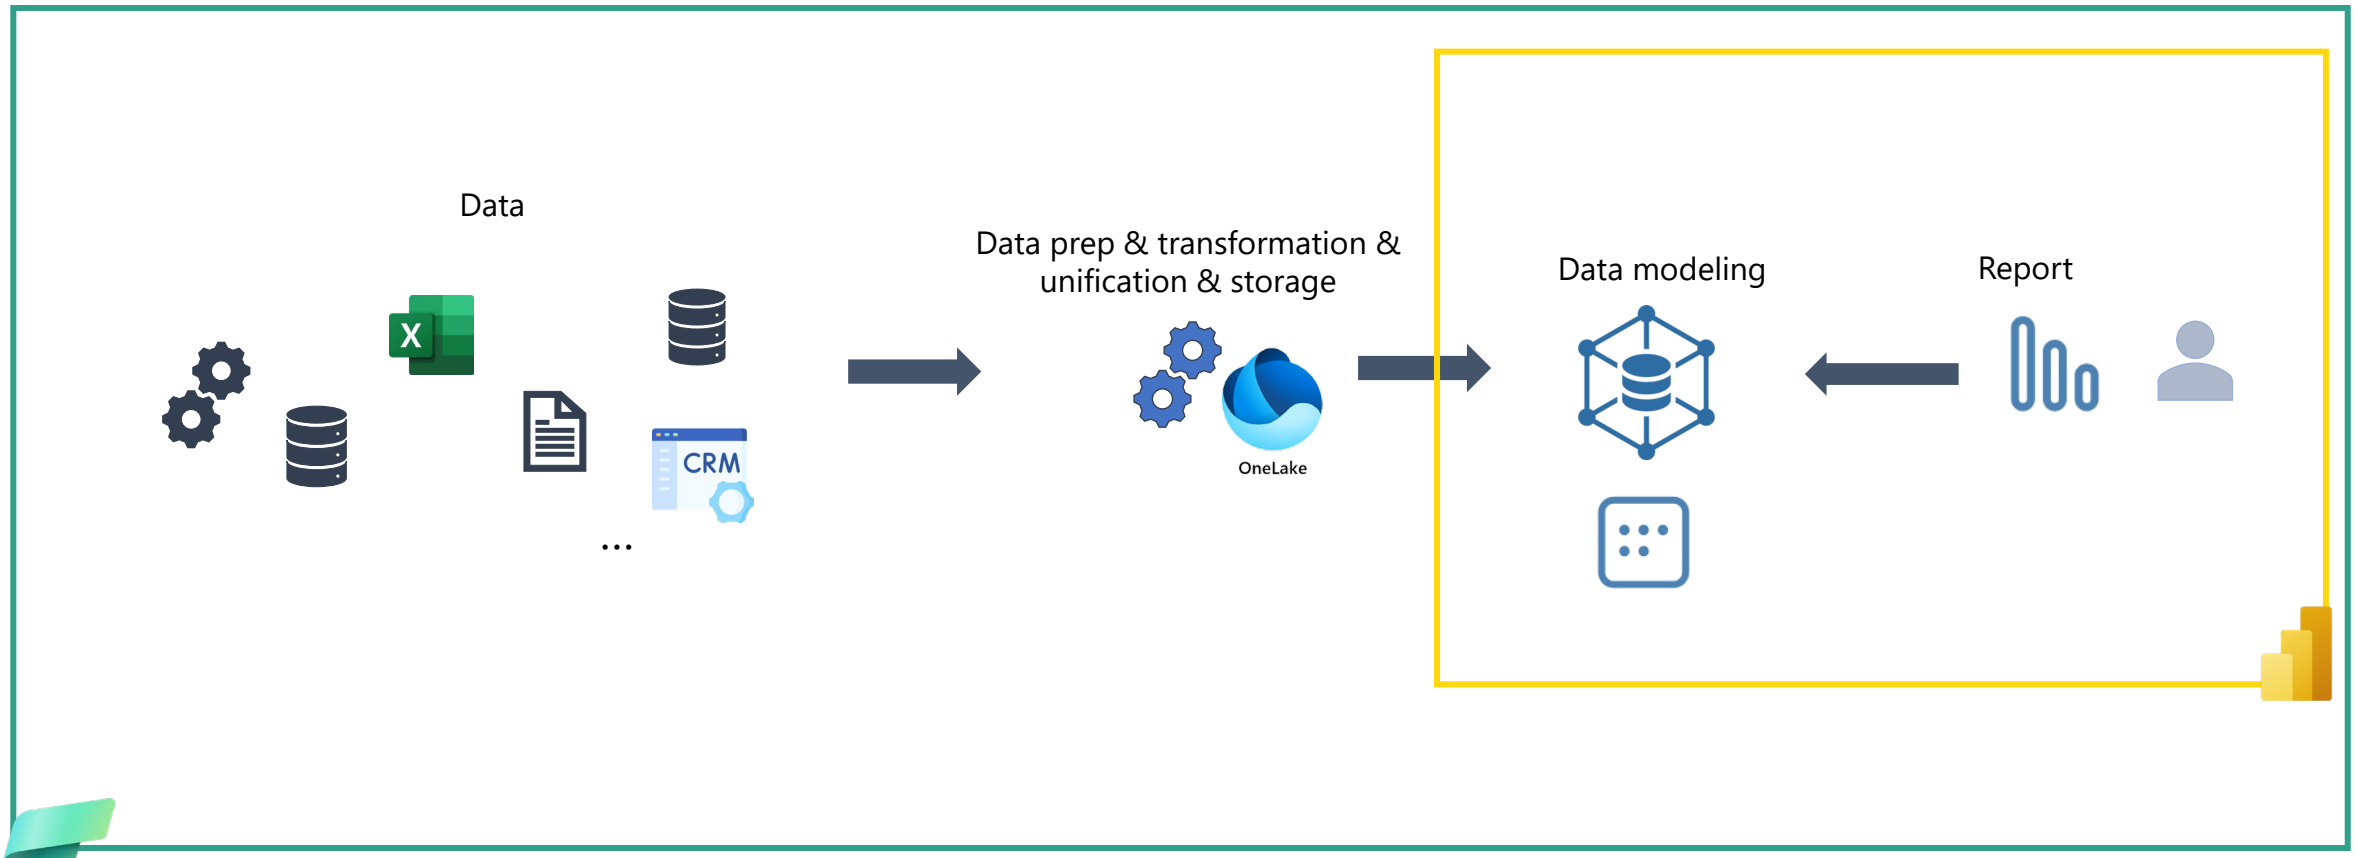
Why this matters for Power BI developers

- One shared, discoverable data foundation
- Less data duplication
- Better visibility into what already exists before building something new

Overall process before Fabric



Overall process with Fabric

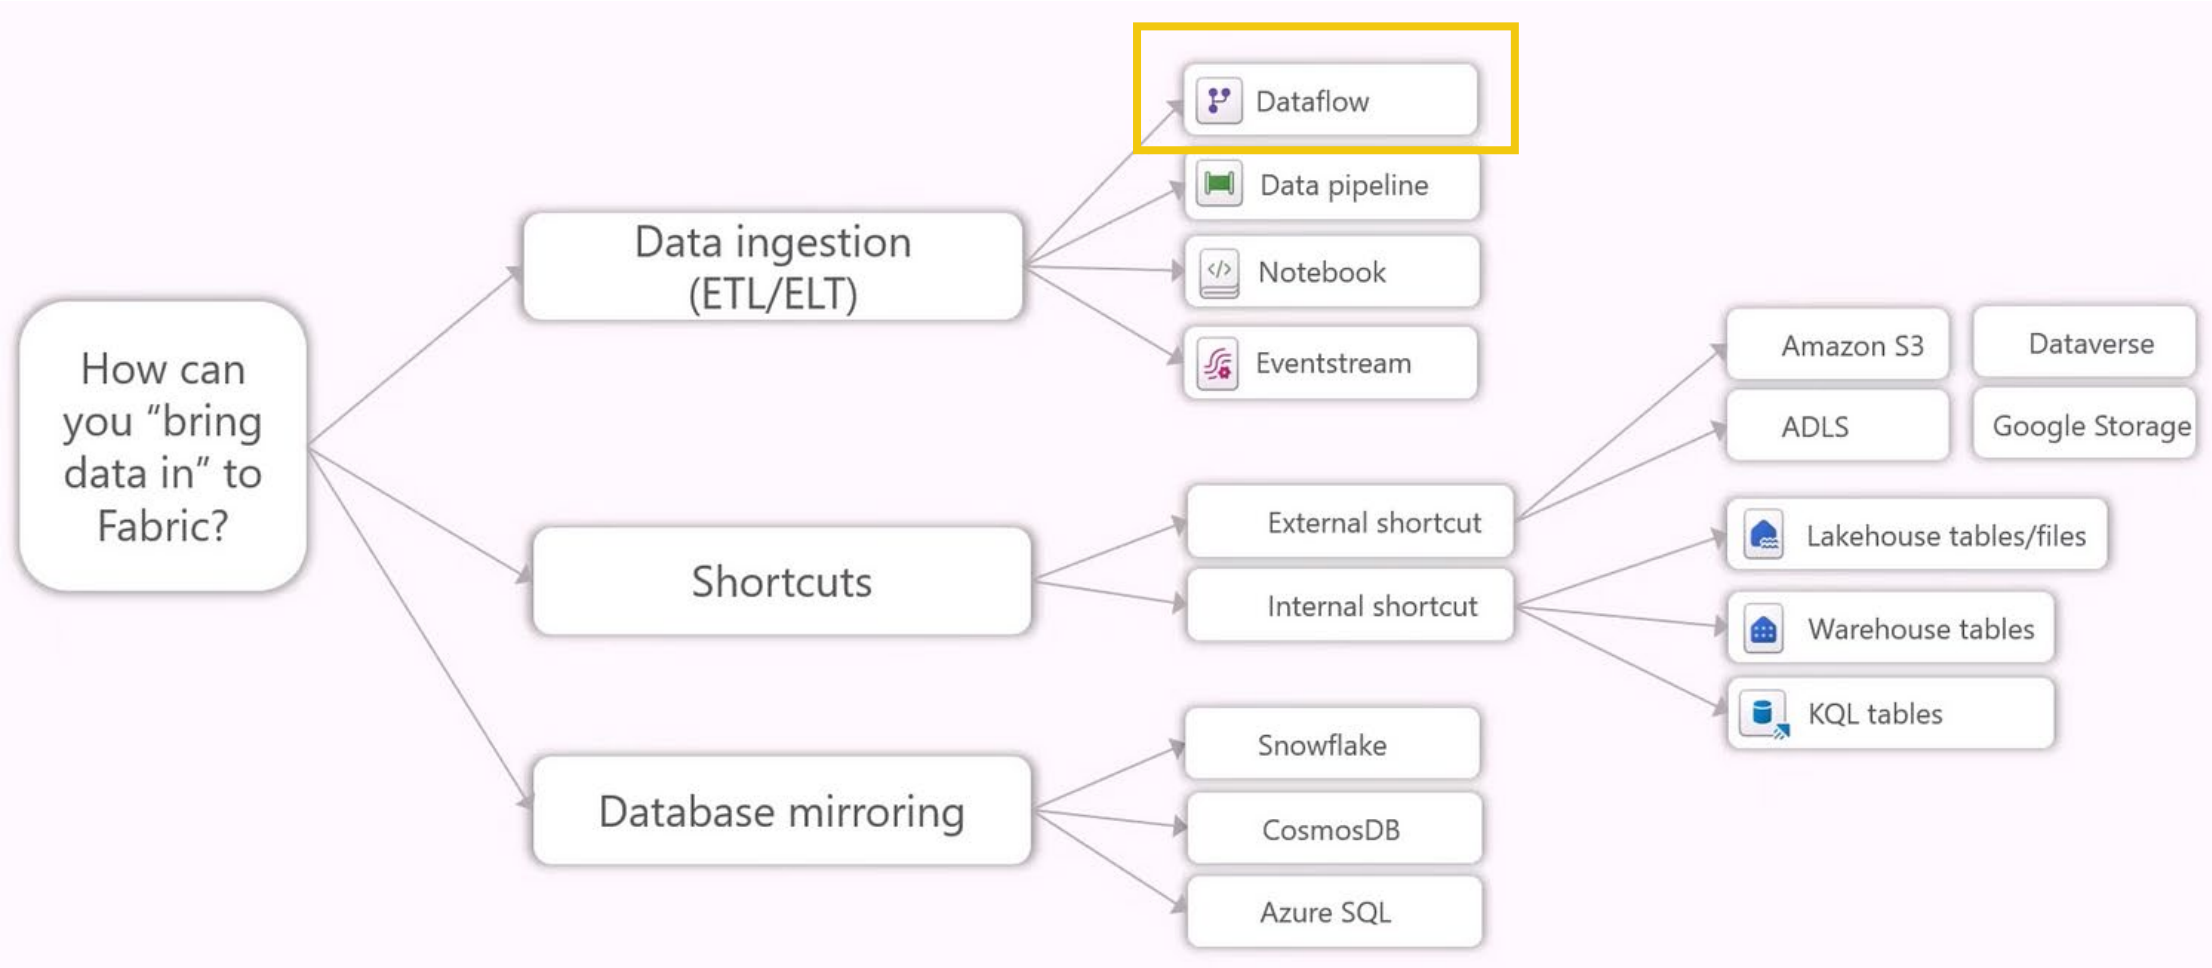


You can as well work in the “old” way

Getting the data



Bringing data into Fabric



Source: "Fabric with Will – 2025"

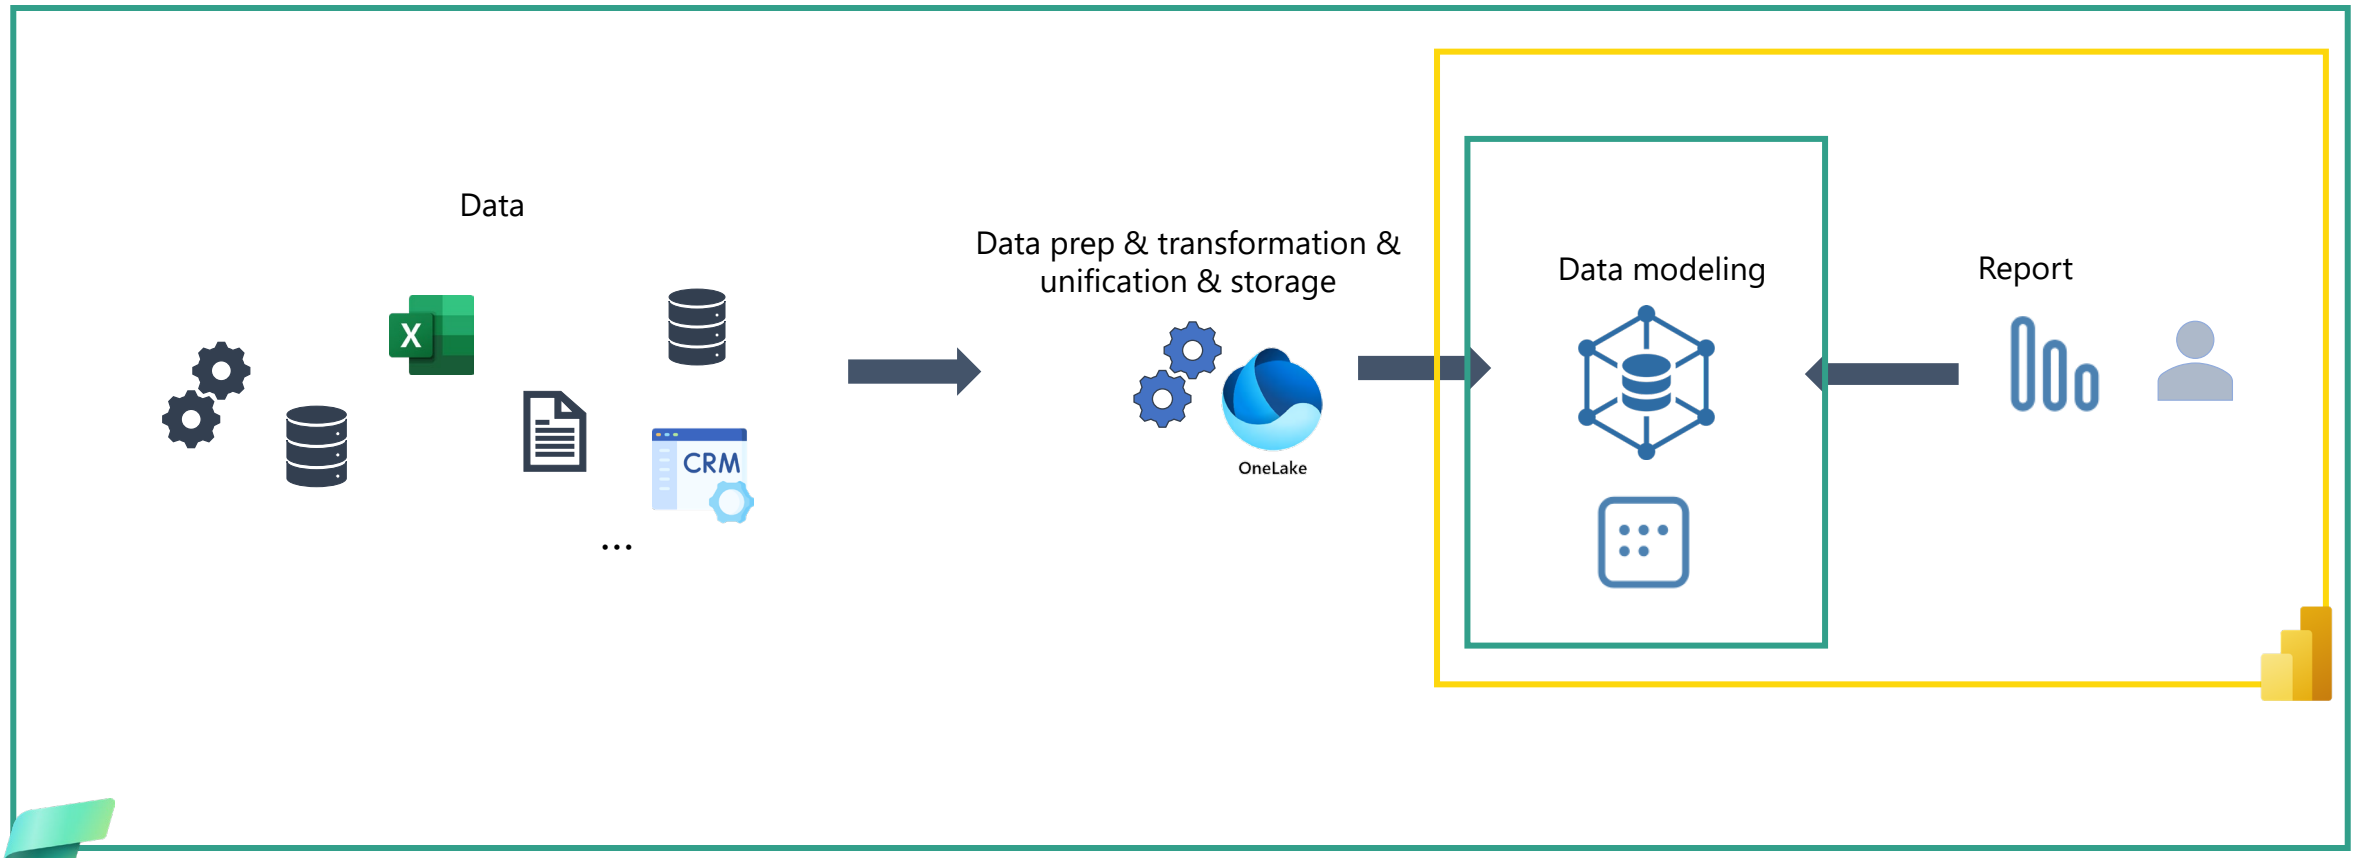
Why this matters for Power BI developers

- More ways to bring data in (low-code, no-code, code)
- Clear separation between data preparation and reporting



Building the model





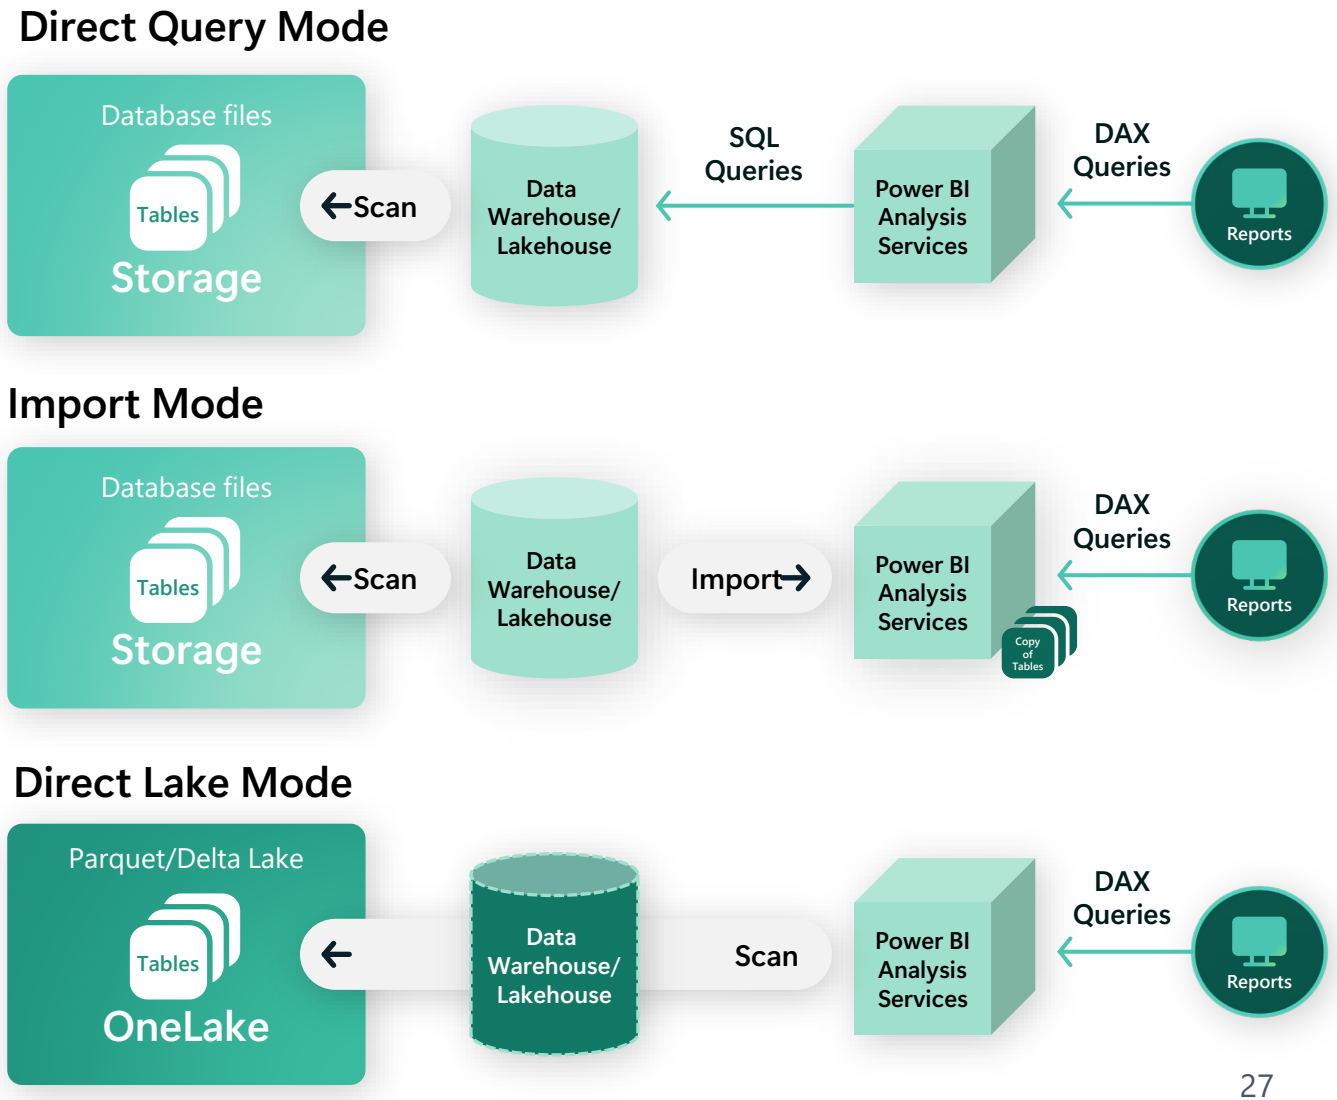
Captain Platform

DirectLake mode

Direct Lake is a fast-path to load the data from the lake straight into the Power BI engine, ready for analysis

Direct Lake is based on loading parquet-formatted files directly from a data lake without having to query a Lakehouse endpoint, and without having to import or duplicate data into a semantic model

Full cloud/web experience!



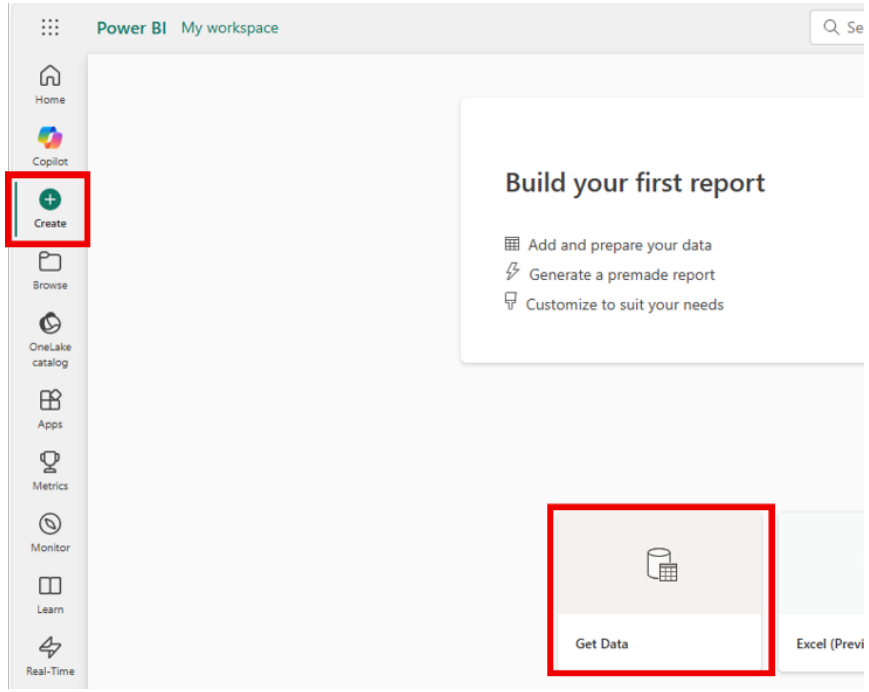
Why this matters for Power BI developers

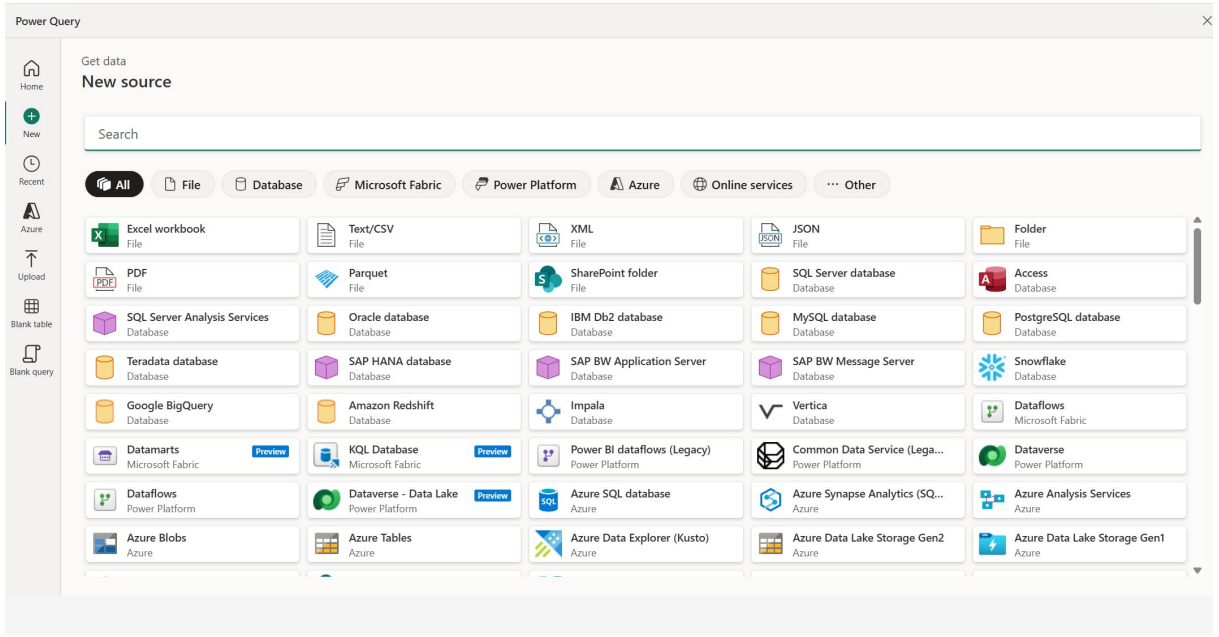
- (Faster) Access to large datasets without import or duplication
- Fewer trade-offs between performance and freshness
- Ability to model directly in the cloud, closer to the data

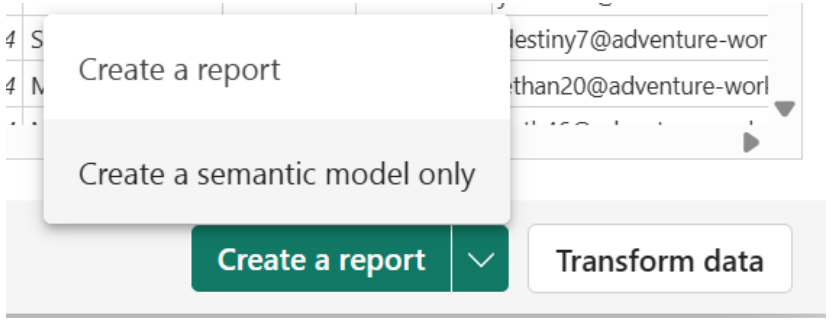
What if I don’t want to connect to data in OneLake?
Can I still create/edit models in Service?



Yes!



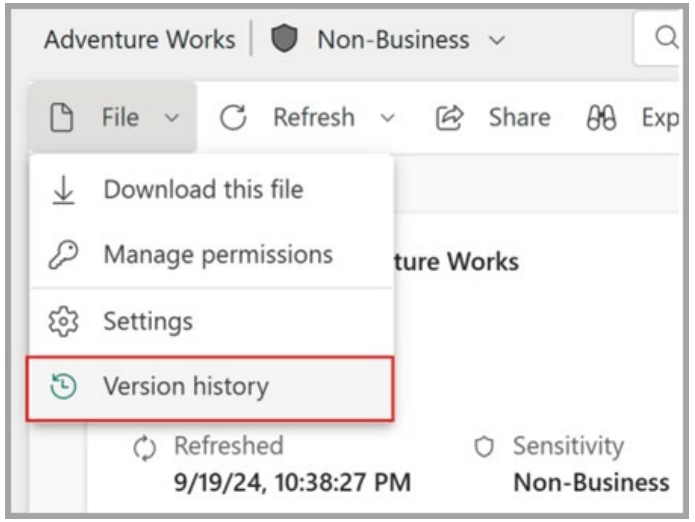


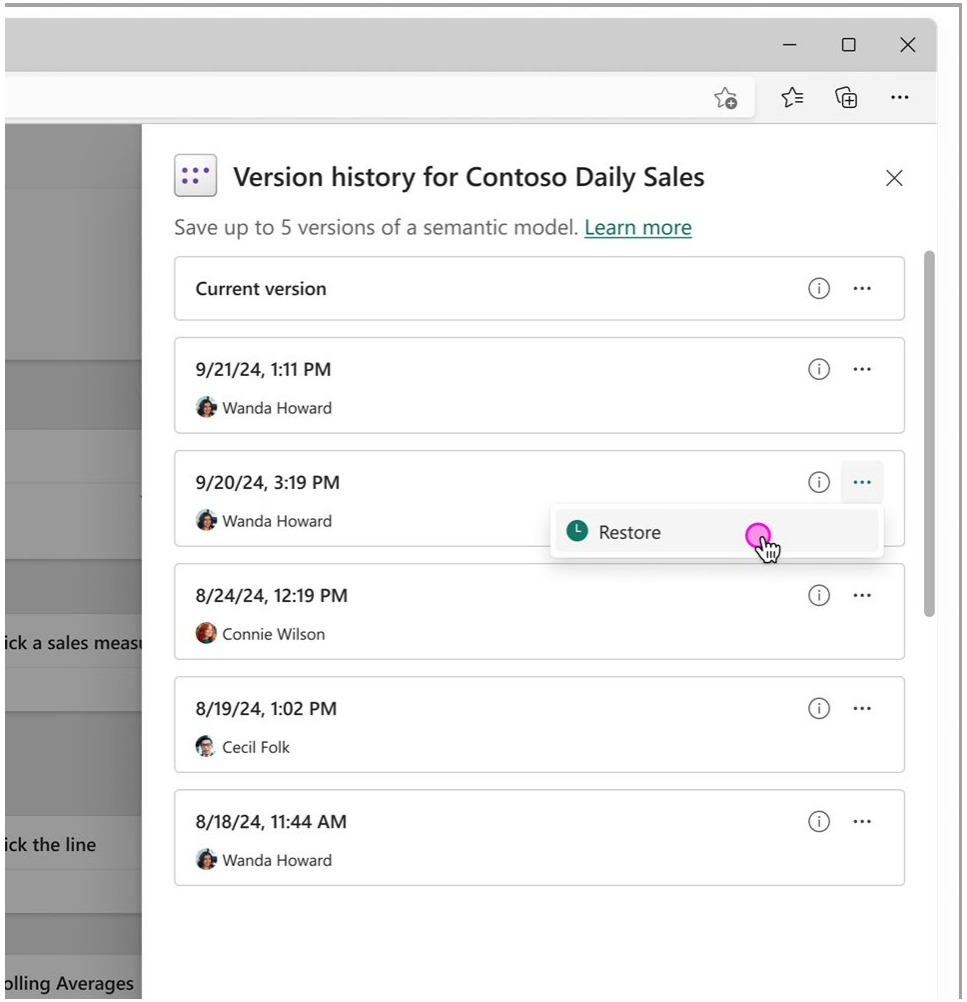


Captain Platform

Web experience - Can I revert if I did a mistake?

As you make changes to your semantic model, your changes will be automatically saved. Semantic model version history is supported for semantic models edited in the web. This feature allows you to recover from critical mistakes.





Have you heard about Vertipaq analyzer?

Or BPA rules?

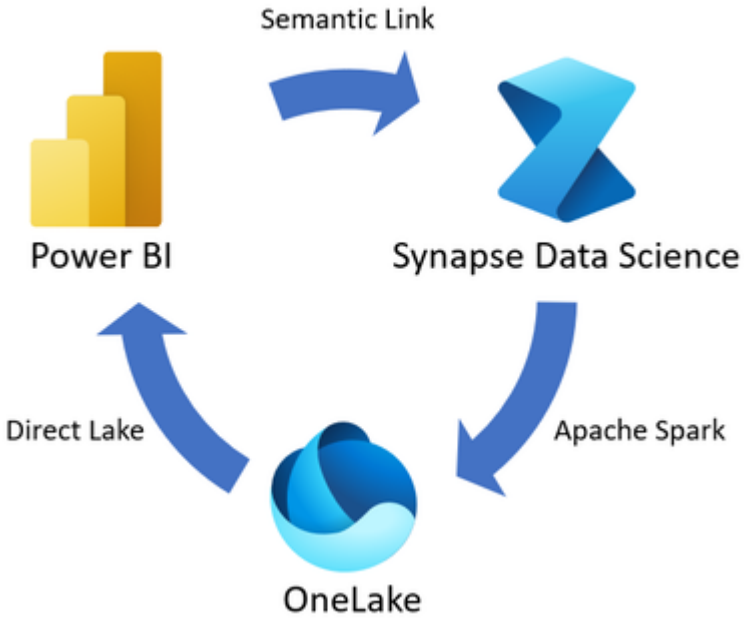
These require the use of external tools*



Semantic Link

Semantic link is a feature that allows you to establish a connection between semantic models and Notebooks in Microsoft Fabric.

- Read data + metadata from Power BI semantic
- Automated documentation
- **Vertipaq Analyzer**
- Translate a semantic model's metadata
- **BPA (Best Practices Analyzer)**
- Rebind reports
- ML using curated semantic model logic.
- Automation/orchestration for refreshes, metadata extraction, or integration into pipelines.



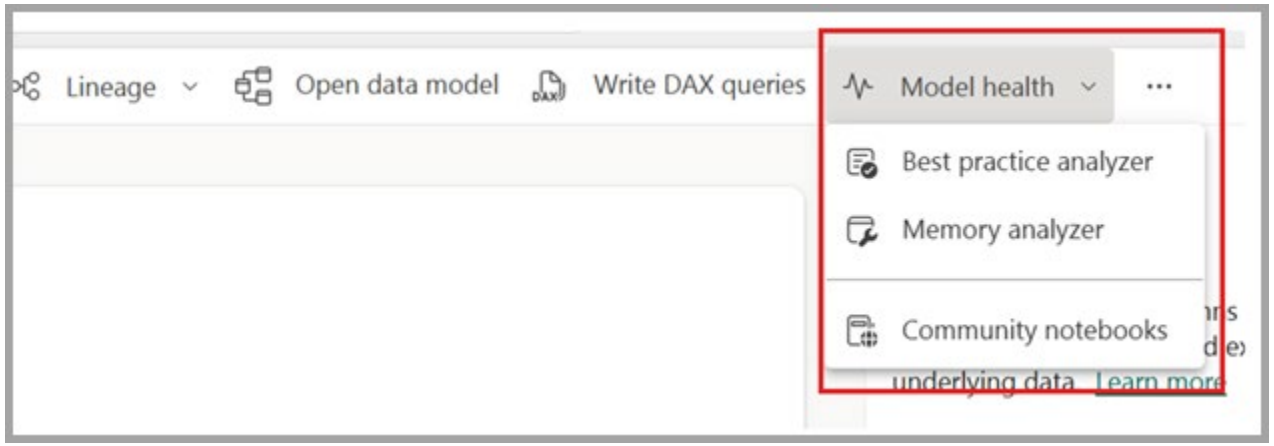
Notebooks? What is that?
Do I need to code something?



No, you can click on a button!

Best Practice Analyzer (BPA) offers tips to improve the design and performance of your semantic model. By default, the BPA checks a set of more than 60 rules against your semantic model and summarizes the results.

The **Memory Analyzer** shows you memory and storage statistics about the objects in your semantic model, such as Tables, Columns, Hierarchies, Partitions, and Relationships.




```
1   fabric.model_memory_analyzer(dataset=dataset, workspace=workspace)
```

41 sec - Command executed in 41 sec 649 ms by Anastasia SALARI on 10:05:18 AM, 4/24/25

Model Summary | Tables | Partitions | Columns (Total Size) | Columns (Temperature) | Relationships | Hierarchies

Columns (Total Size)

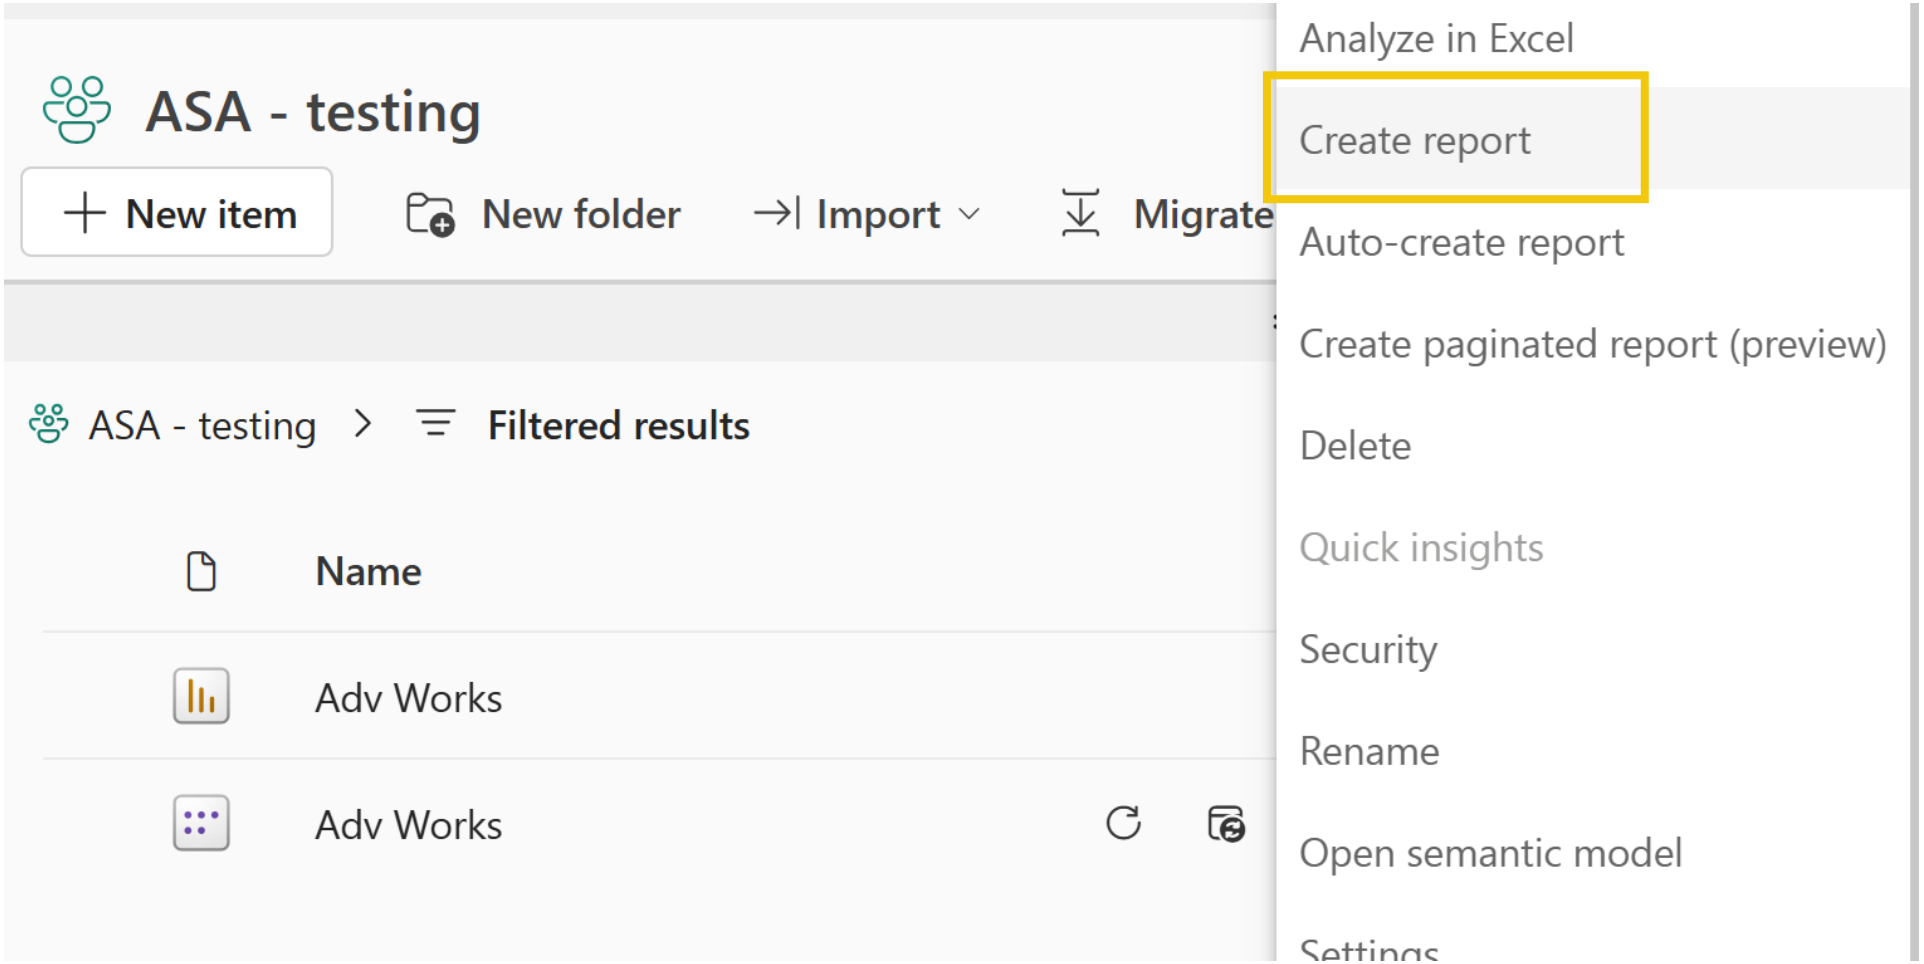
	Table Name	Column Name	Type	Data Type	Cardinality	Total Size	Data Size	Dictionary Size	Hierarchy Size	% Table	% DB	Encoding	Is Resident	Temperat
0	Complains	Consumer complaint narrative	Data	String	454,858	334.95 MB	1.22 MB	330.26 MB	3.47 MB	93.13%	93.06%	Hash	True	
1	Complains	Complaint ID	Data	Int64	1,471,766	11.23 MB	5.61 MB	128 B	5.61 MB	3.12%	3.12%	Value	True	
2	Complains	ZIP code	Data	String	23,546	3.69 MB	2.56 MB	973.59 KB	184.00 KB	1.03%	1.03%	Hash	True	
3	Complains	Date sent to company	Data	DateTime	2,908	2.36 MB	2.24 MB	104.12 KB	22.77 KB	0.66%	0.66%	Hash	True	
4	Complains	Date received	Data	String	2,959	2.36 MB	2.24 MB	99.96 KB	23.16 KB	0.66%	0.66%	Hash	True	
5	Complains	Company	Data	String	5,497	1.39 MB	1.11 MB	242.50 KB	42.98 KB	0.39%	0.39%	Hash	True	
6	Complains	Col len	Calculated	Int64	8,407	1.14 MB	937.37 KB	161.31 KB	65.72 KB	0.32%	0.32%	Hash	True	

Captain Platfo

Building reports



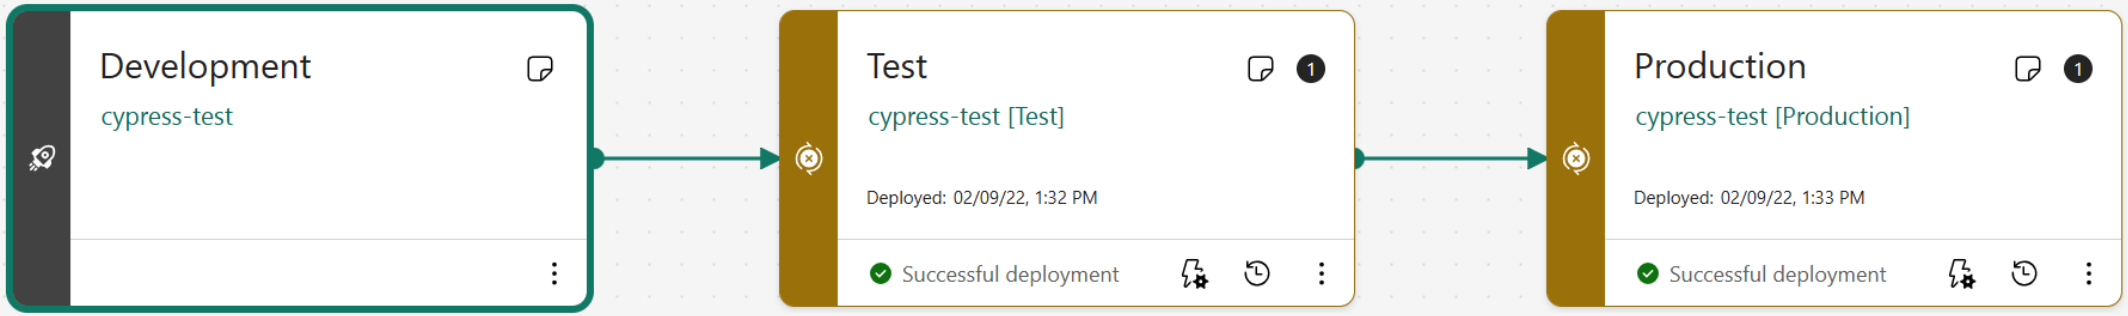
Building reports in Power BI Service



Overall (dev) workflow



Deployment pipelines (a Fabric feature)



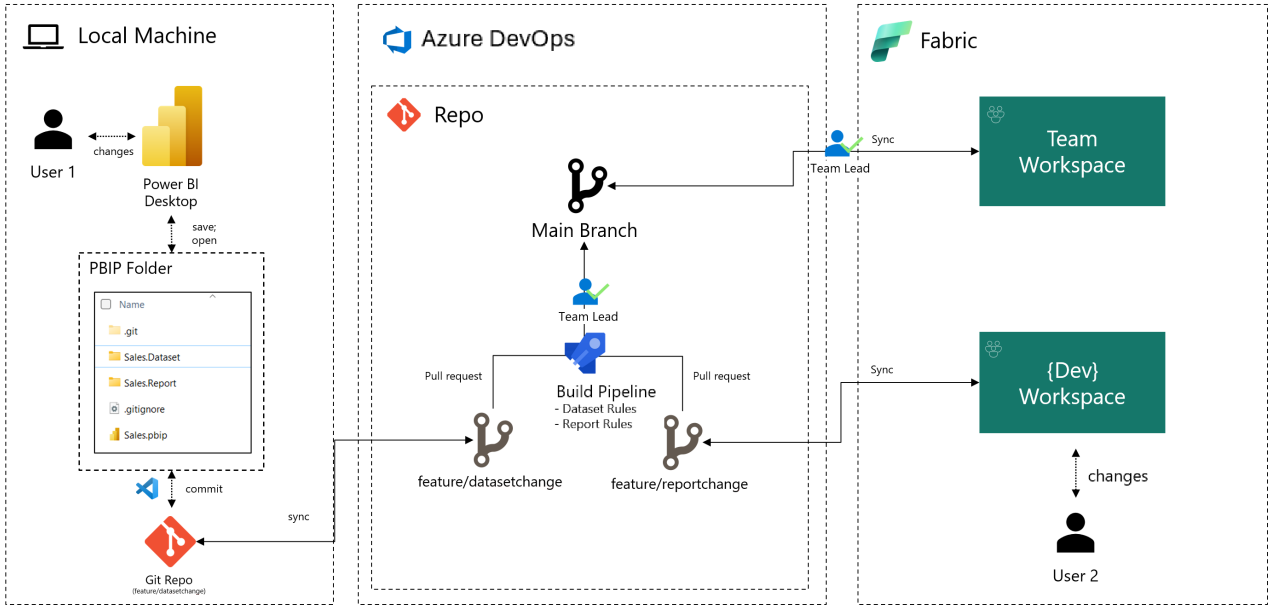
Captain Platform

Git & DevOps integration in Fabric

- Native Git integration in Microsoft Fabric (Azure DevOps & GitHub)
- Store and version Power BI semantic models, reports, and Fabric items as code
- Track who changed what, when, and why through commit history
- Enable code reviews, pull requests, and validation before deployment
- Support CI/CD workflows to promote changes safely across environments

Why it matters for Power BI developers

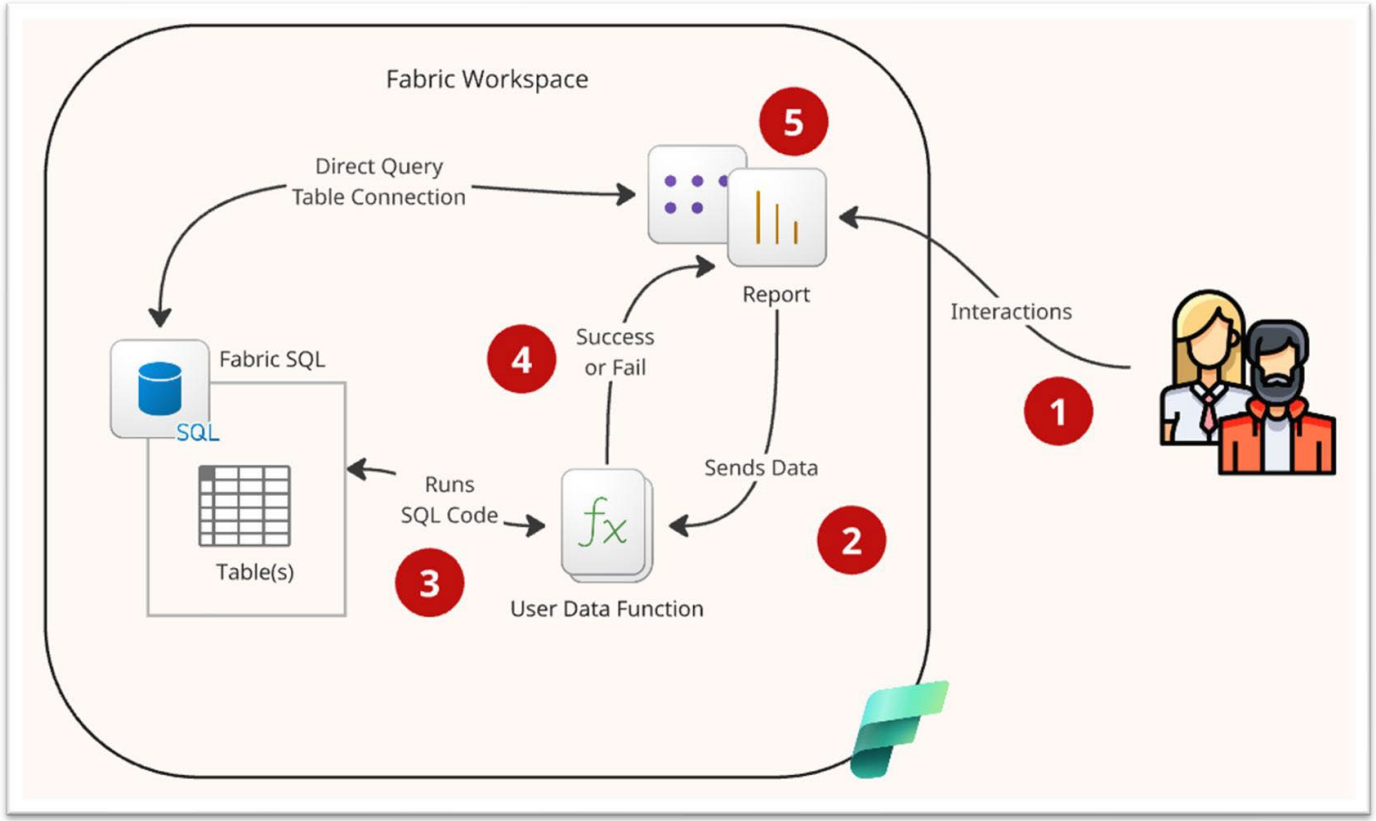
- Better collaboration in multi-developer teams
- Fewer manual deployments and fewer production mistakes

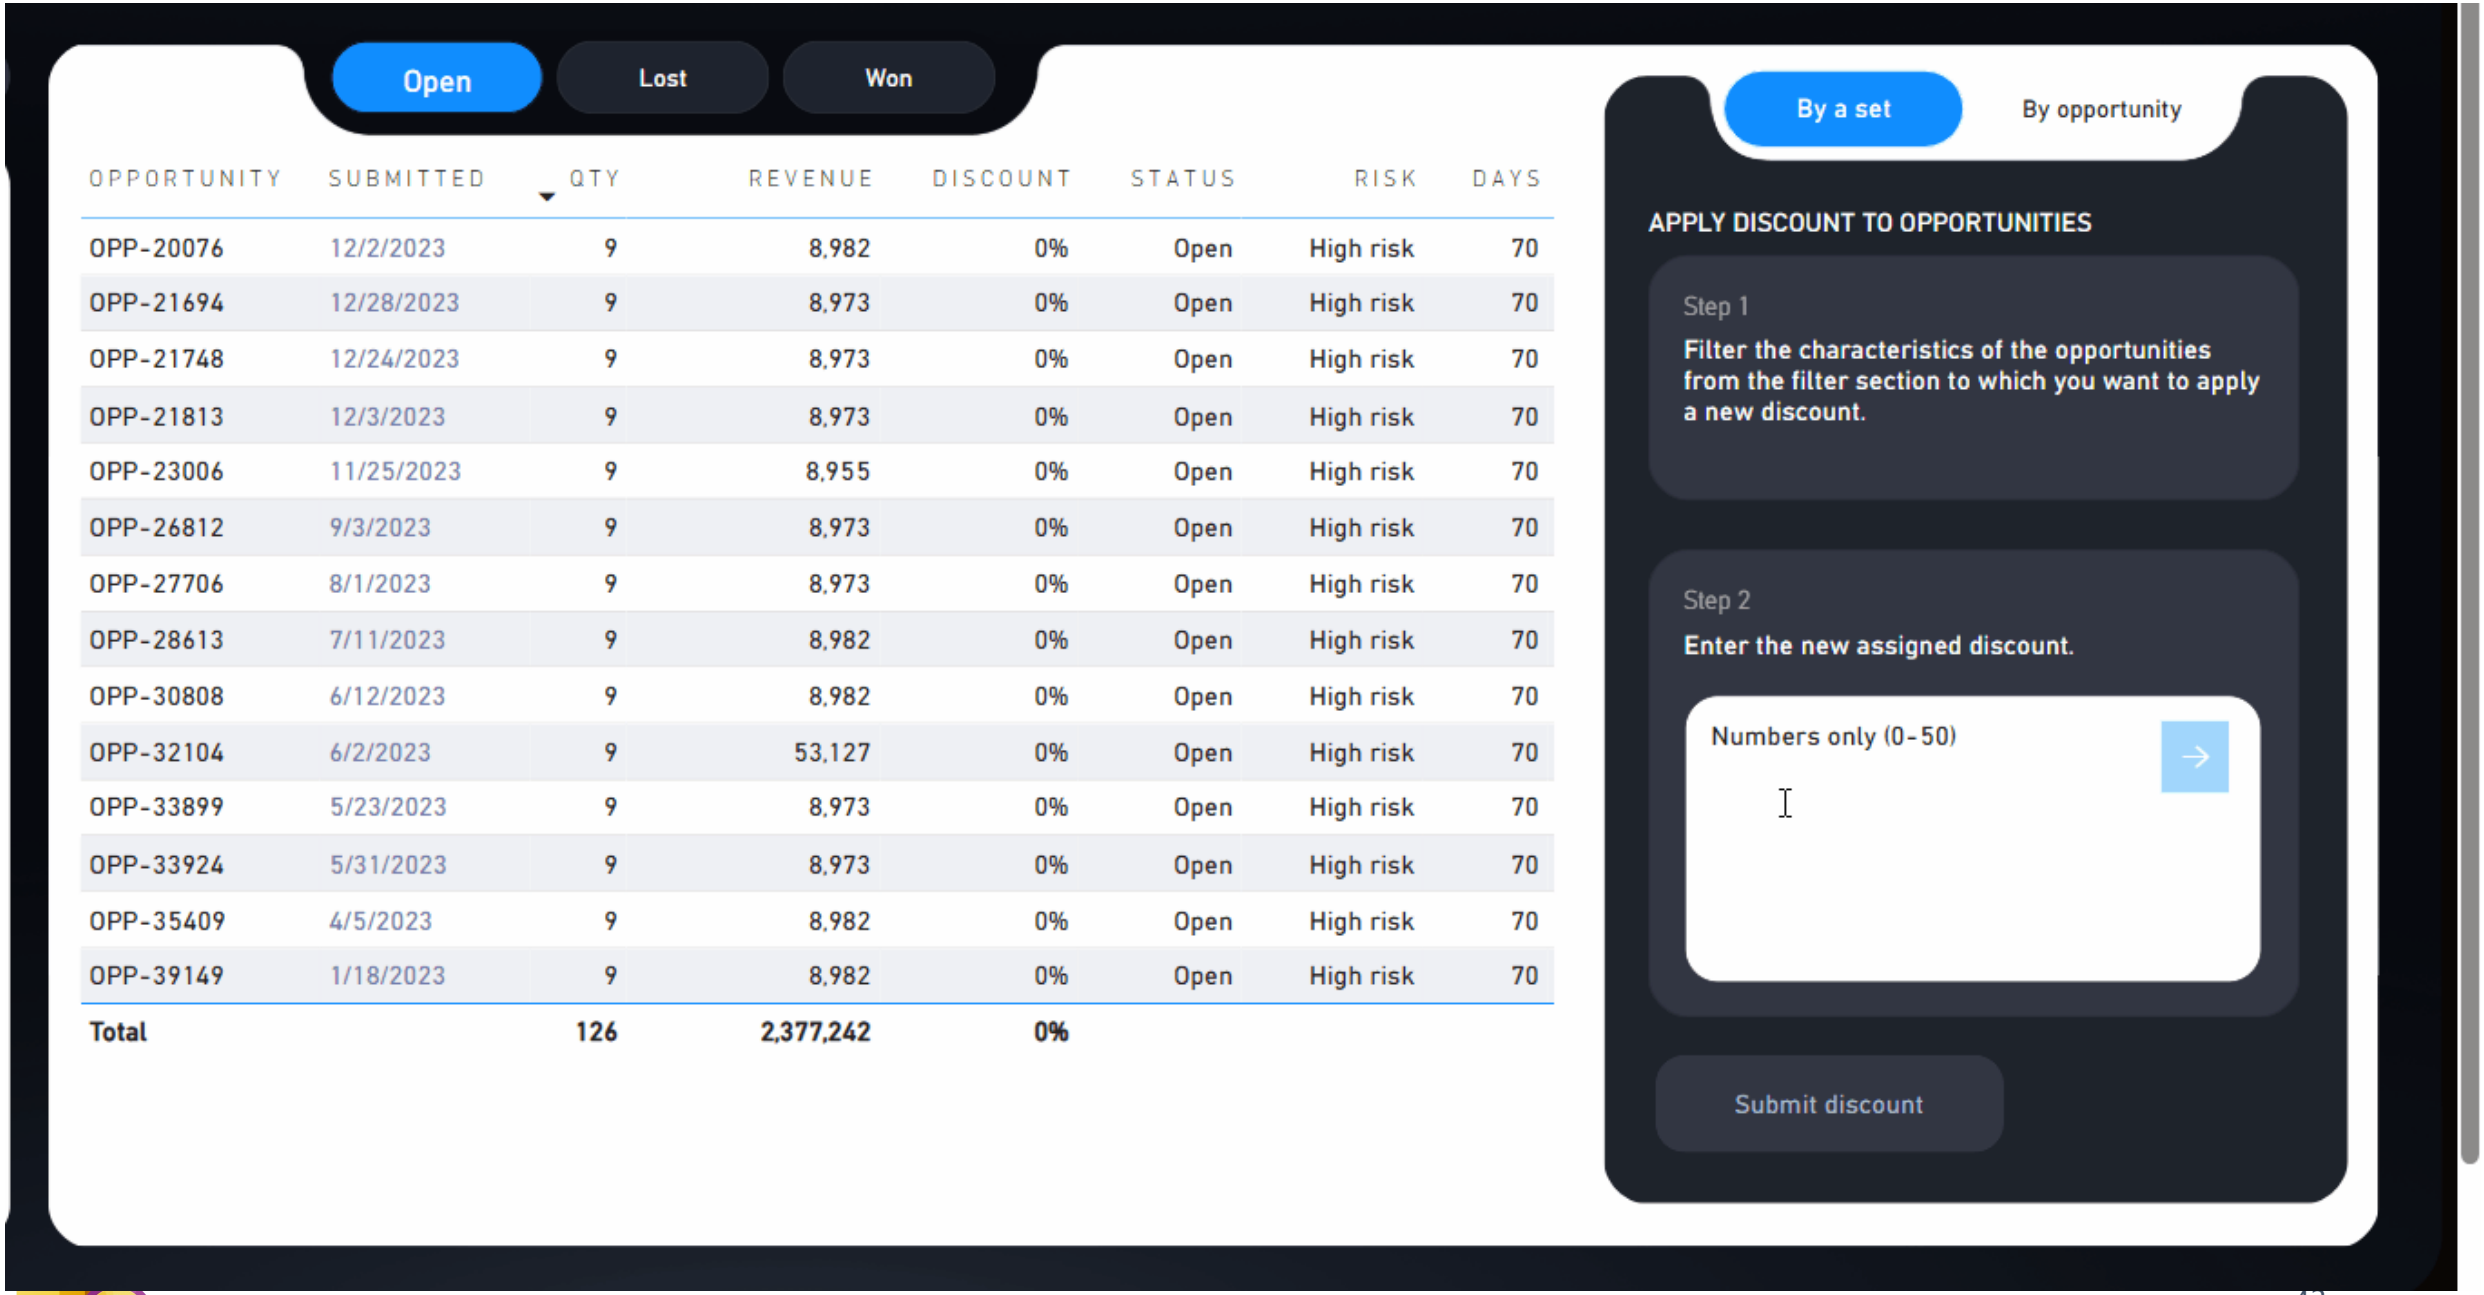


New integrations

Translytical task flows

With translytical task flows, you can automate end-user actions like updating records, adding annotations, or creating workflows that trigger actions in other systems.




Open
Lost
Won
OPPORTUNITY SUBMITTED QTY REVENUE DISCOUNT STATUS RISK DAYS
OPP-20076 12/2/2023 9 8,982 0% Open High risk 70
OPP-21694 12/28/2023 9 8,973 0% Open High risk 70
OPP-21748 12/24/2023 9 8,973 0% Open High risk 70
OPP-21813 12/3/2023 9 8,973 0% Open High risk 70
OPP-23006 11/25/2023 9 8,955 0% Open High risk 70
OPP-26812 9/3/2023 9 8,973 0% Open High risk 70
OPP-27706 8/1/2023 9 8,973 0% Open High risk 70
OPP-28613 7/11/2023 9 8,982 0% Open High risk 70
OPP-30808 6/12/2023 9 8,982 0% Open High risk 70
OPP-32104 6/2/2023 9 53,127 0% Open High risk 70
OPP-33899 5/23/2023 9 8,973 0% Open High risk 70
OPP-33924 5/31/2023 9 8,973 0% Open High risk 70
OPP-35409 4/5/2023 9 8,982 0% Open High risk 70
OPP-39149 1/18/2023 9 8,982 0% Open High risk 70
Total 126 2,377,242 0%
By a set
By opportunity
APPLY DISCOUNT TO OPPORTUNITIES
Step 1
Filter the characteristics of the opportunities from the filter section to which you want to apply a new discount.
Step 2
Enter the new assigned discount.
Numbers only (0-50)
Submit discount

Data activator

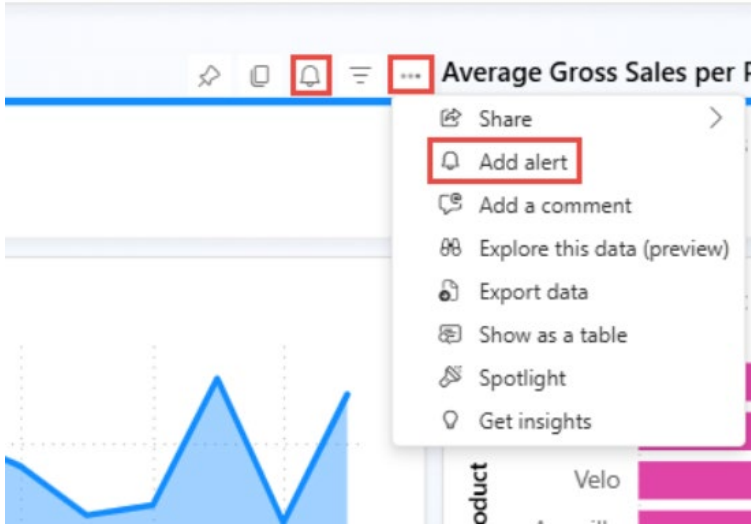
A no-code experience in Microsoft Fabric for automatically taking actions when patterns or conditions are detected in changing data.

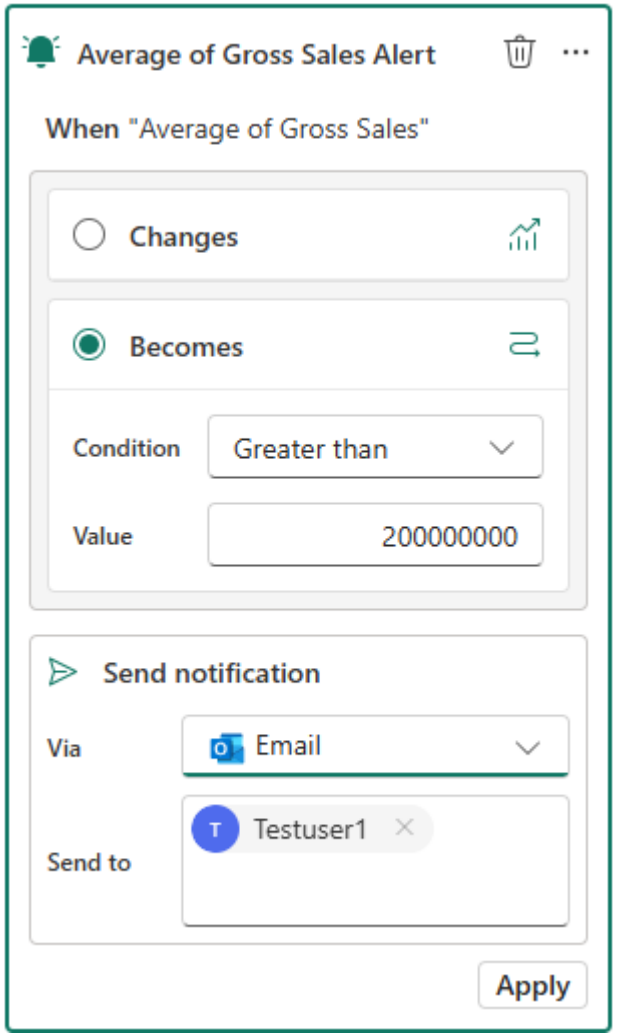
Automatically monitor Power BI semantic models.

Define no-code rules to detect actionable conditions in your data.

Send alerts in Microsoft Teams or email when actionable conditions occur.

Combine Data Activator with Power Automate to send alerts via 3rd party systems, or trigger actions in ticketing systems and other apps.





Data agents

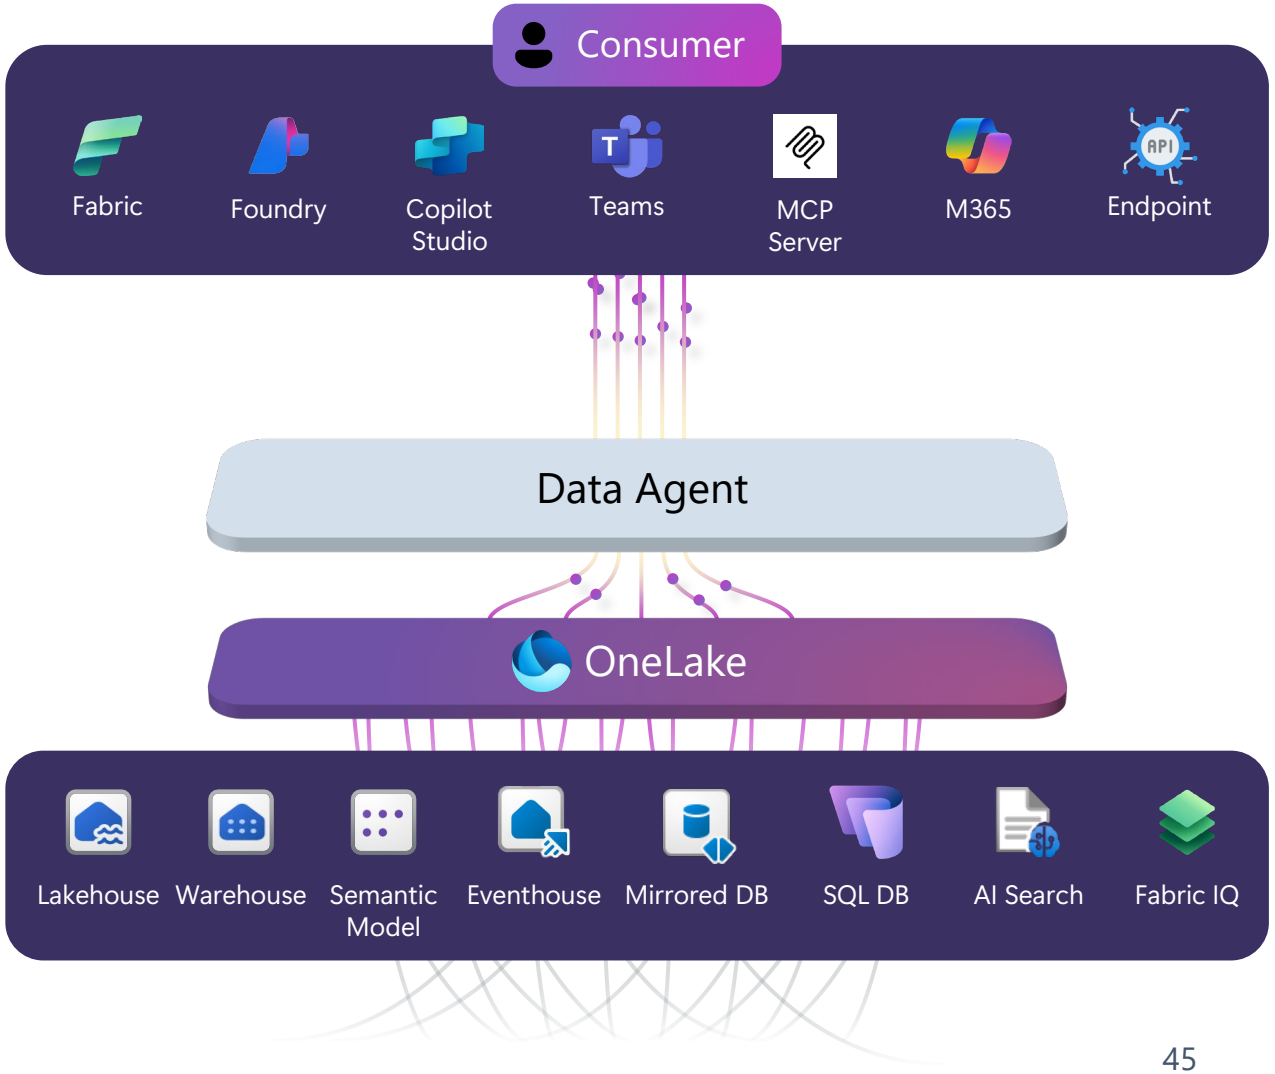
Fabric Data Agents are AI analysts that let users talk to their data in Microsoft Fabric.

Understand natural language questions
Translate them into **SQL, KQL, or DAX** depending on data source

Query Lakehouses, Warehouses, SQL DBs, KQL DBs, Semantic Models

Return contextualized insights

Use instructions, examples, and tools to produce reliable answers

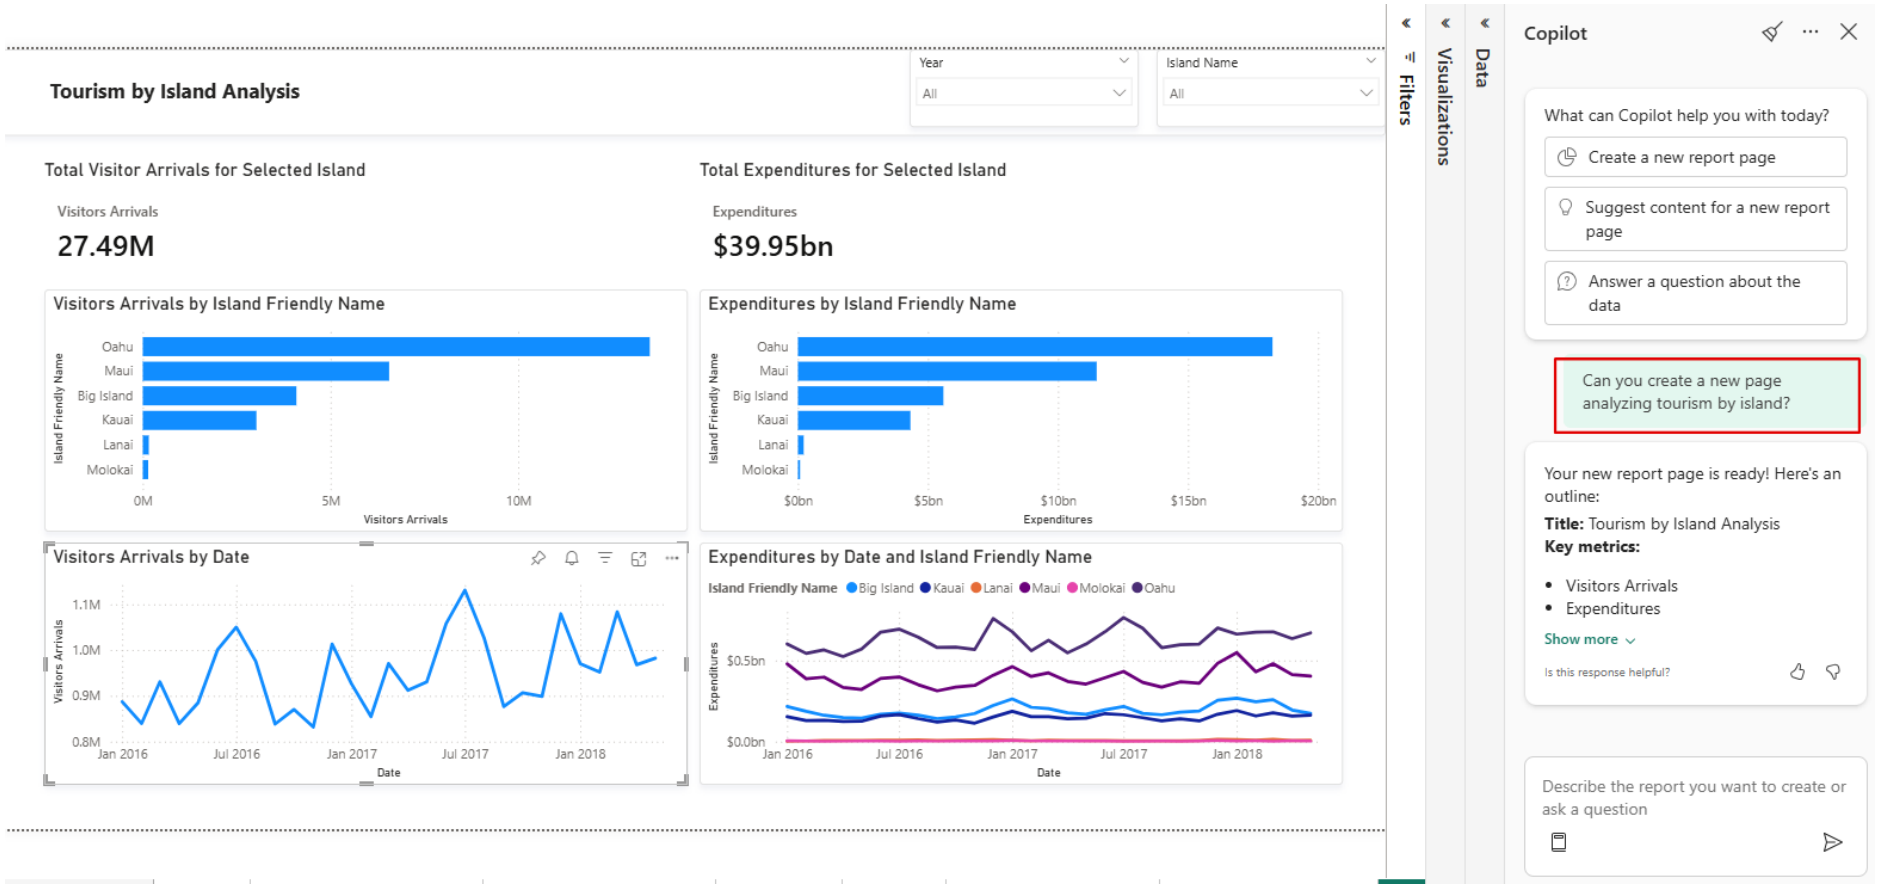


Captain Platform

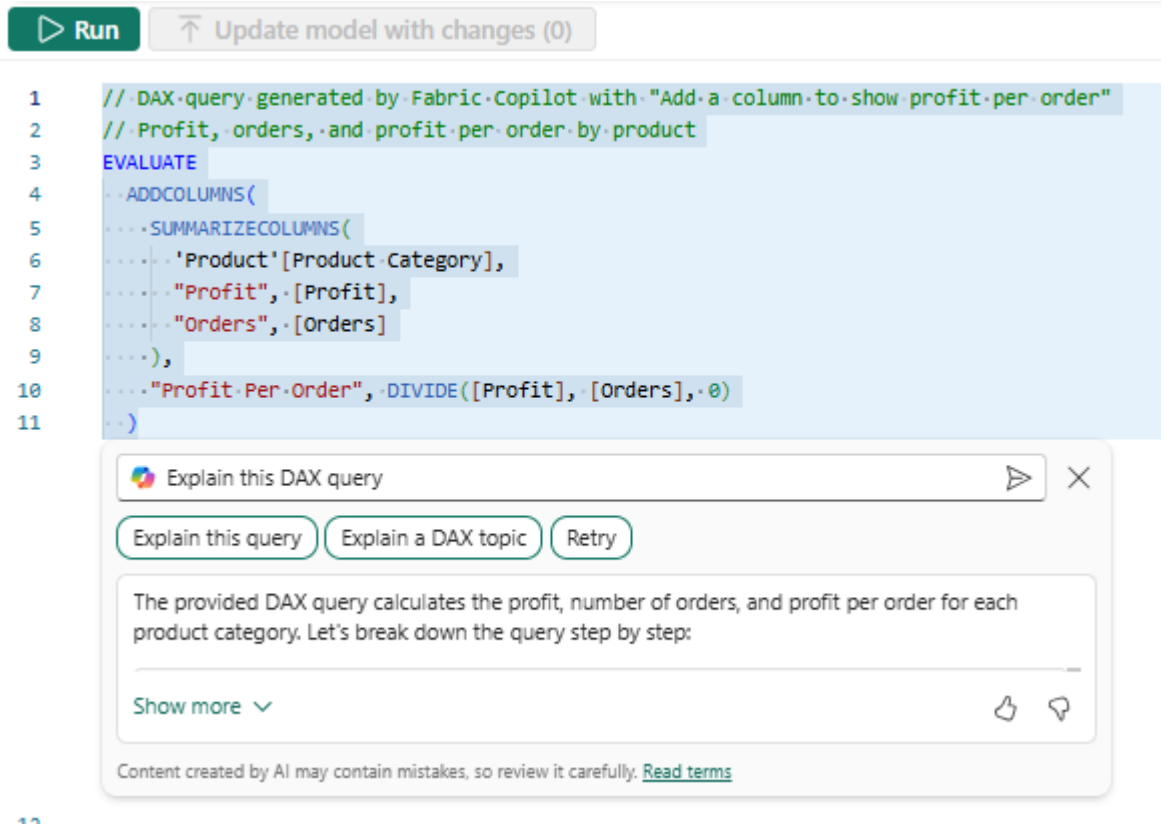
Integration with Copilot

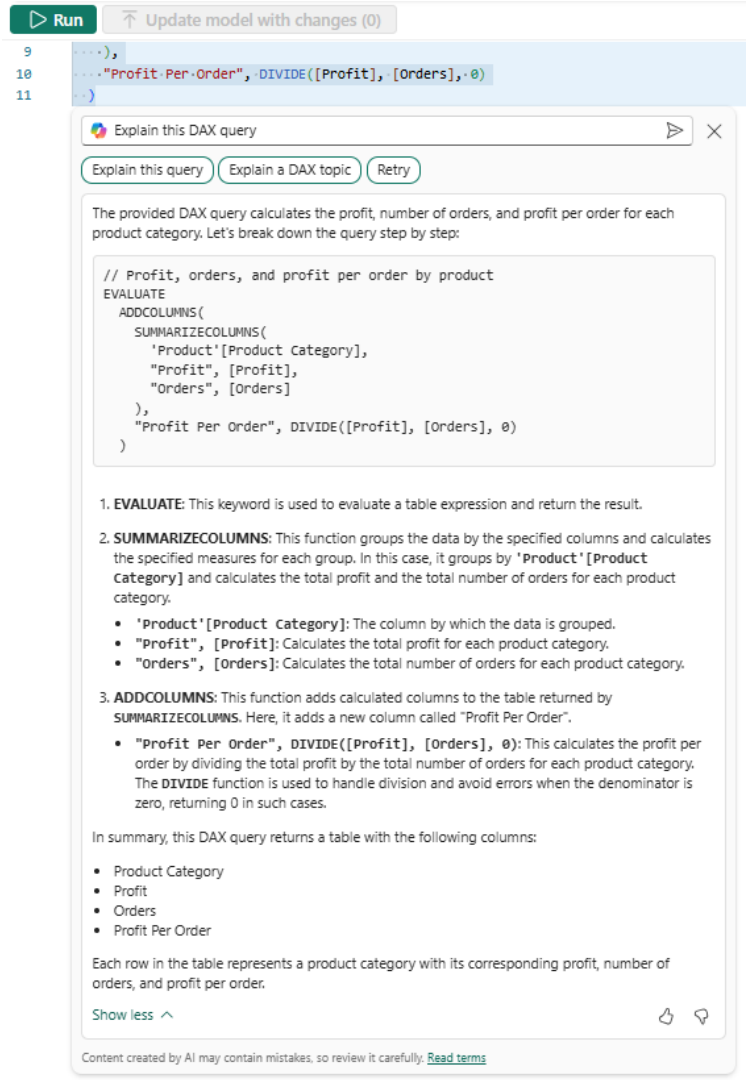


Copilot for report build/edit



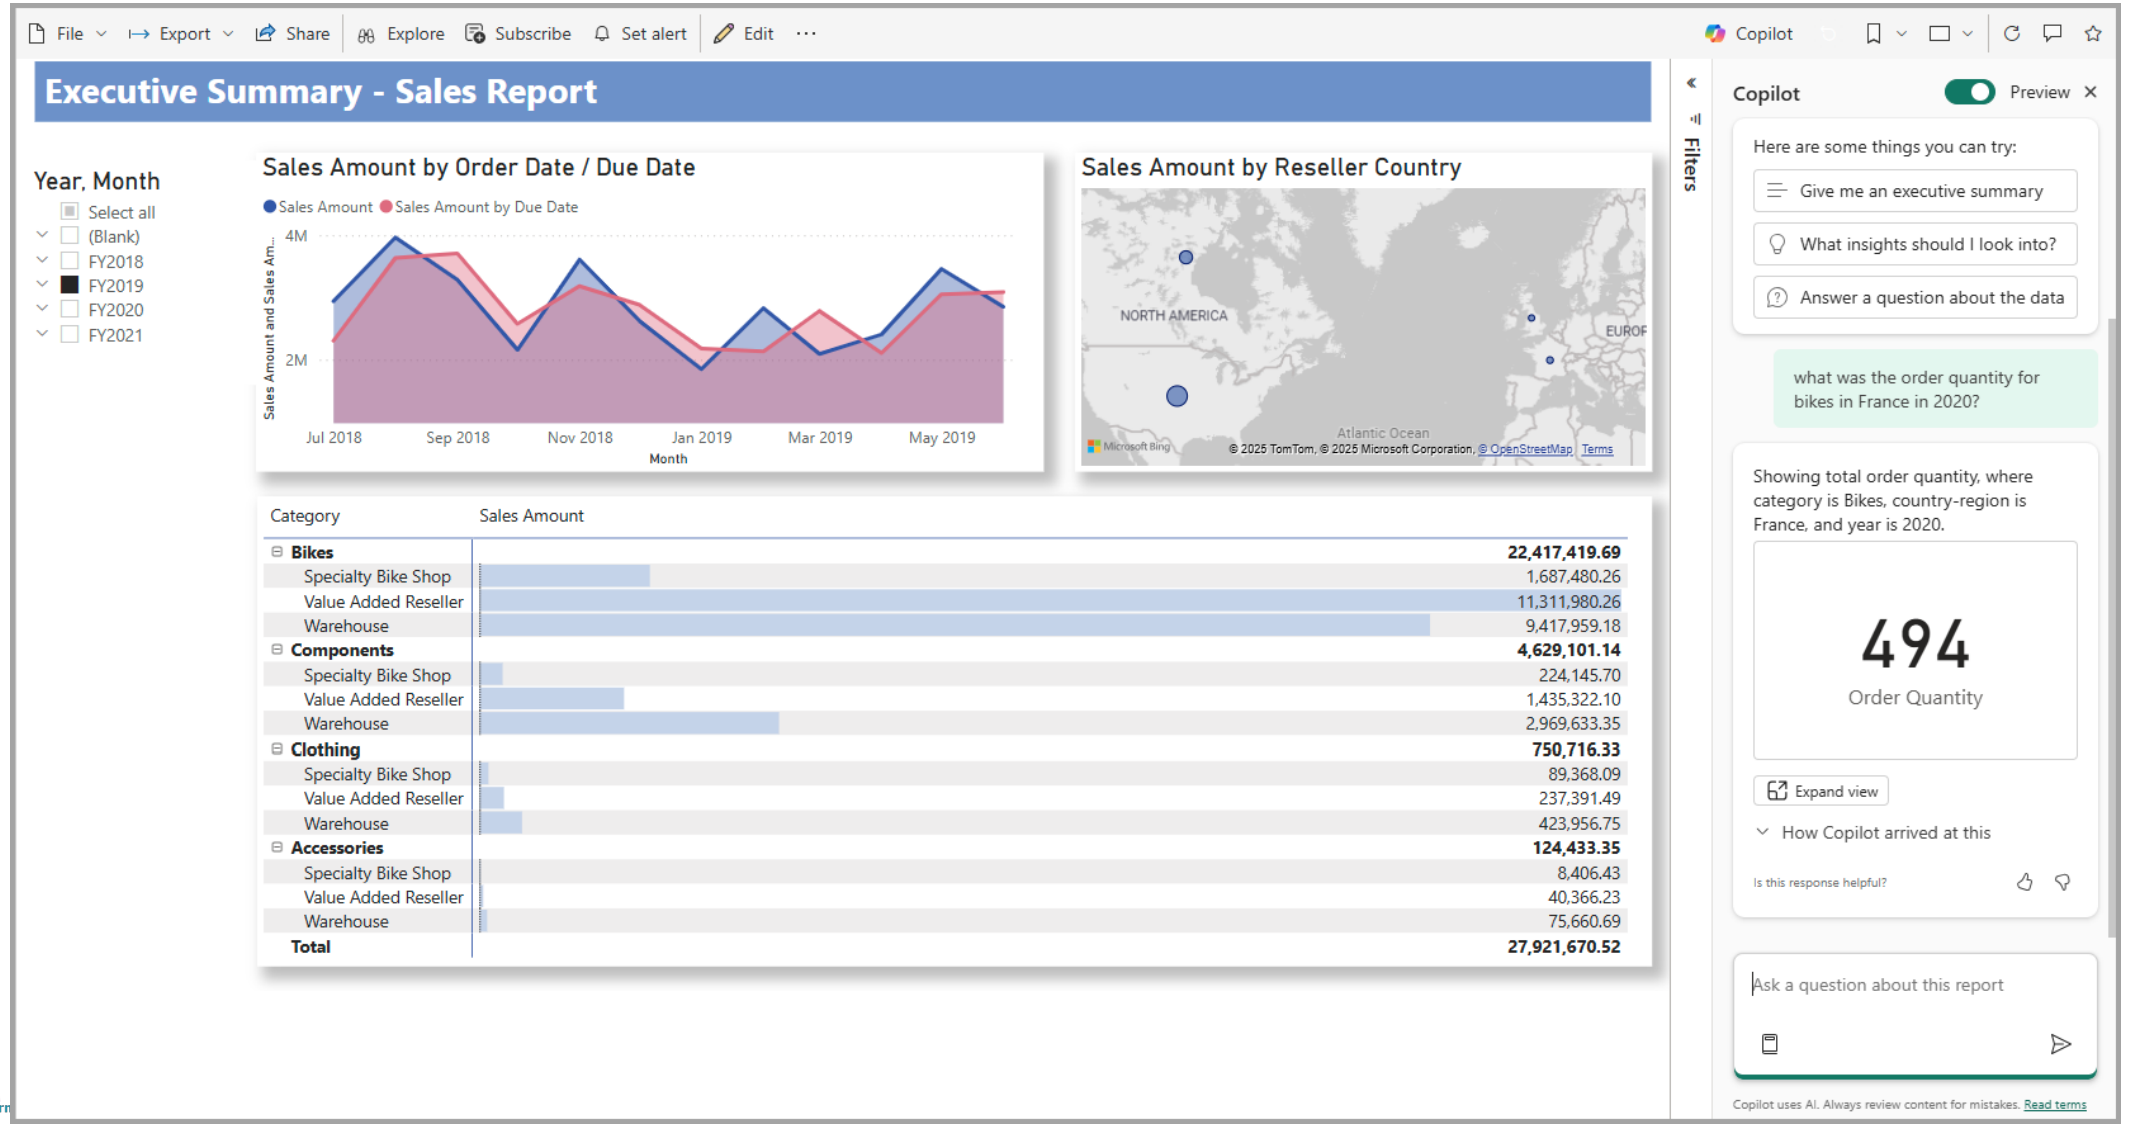
Copilot for Power BI - DAX





Captain Platform

Copilot “Ask questions on your data”

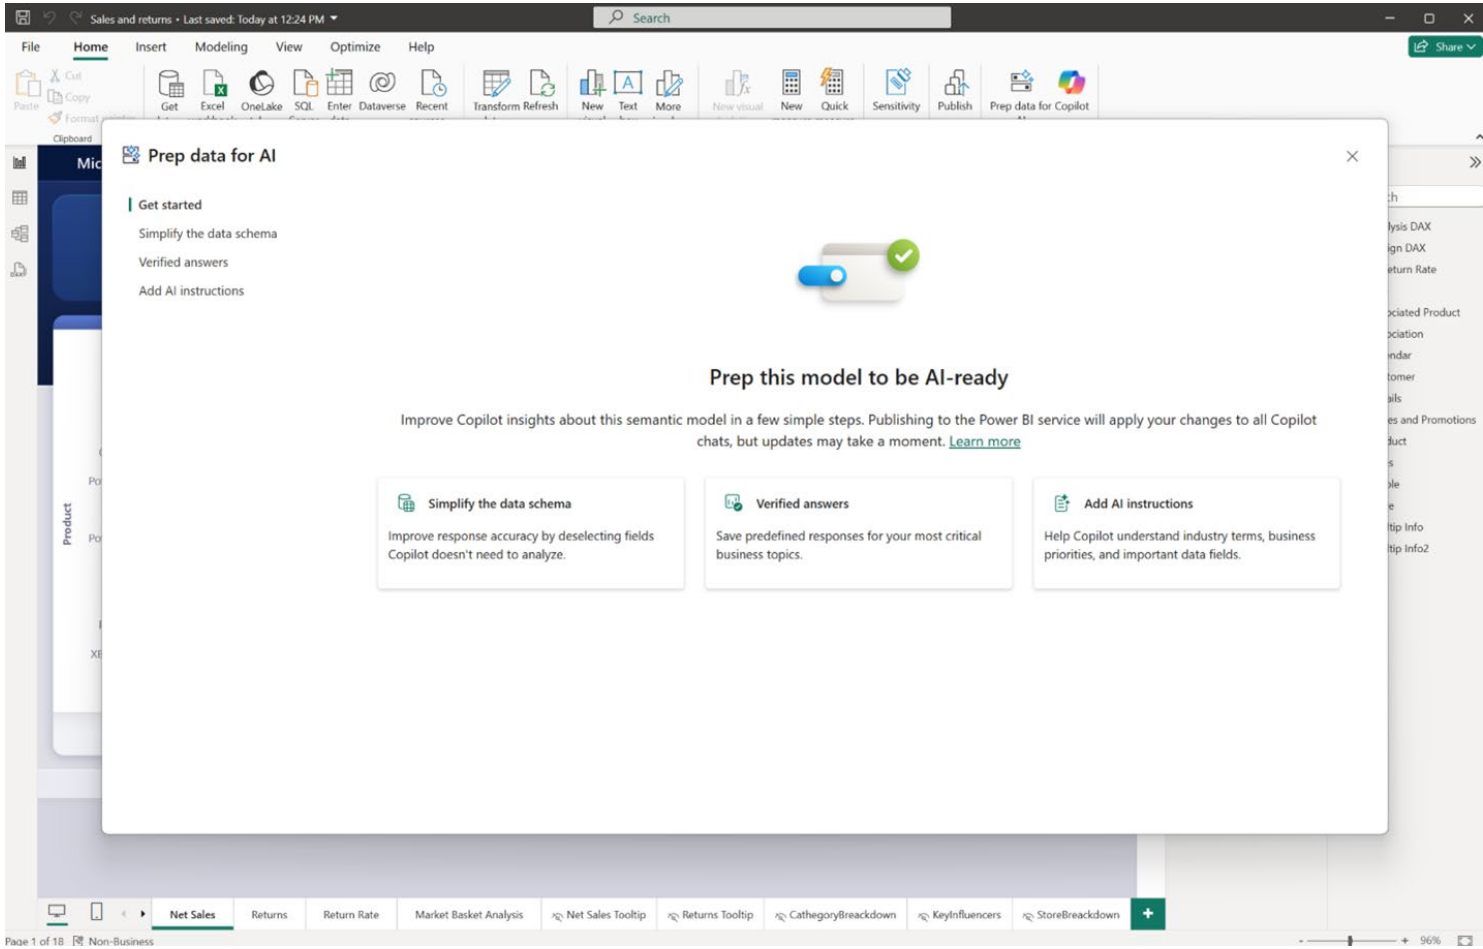


New responsibilities related to AI



Prep data for AI

- Preparing your data improves Copilot accuracy and relevance
- Power BI introduced 3 new “Prep data for AI” features
  - These features let you define model context, add AI instructions, and guide Copilot toward the right fields and tables.
- Verified answers help control and standardize Copilot outputs
- You can mark a semantic model as “Approved for Copilot”



Content types Copilot looks at

Copilot can search through content metadata to find the most relevant item.

The following list includes the information within an item that Copilot looks at:

For reports:

- Page names
- Visual titles
- Filter pane titles
- Text boxes

For semantic models:

- Table names
- Table descriptions
- Column names
- Column descriptions
- Measure names
- Measure descriptions
- Verified answer prompts
- Verified answer charts



Semantic models become central assets reused across workloads, not just report backends.

All these things bring also more complex security & governance



Power BI x Fabric

What changes
- Power BI is no longer isolated, it's part of a broader data platform
- Data, models, and workflows are more connected, discoverable, and reusable
- New capabilities (DirectLake, CI/CD, AI) expand what's possible

What doesn't
- Good data modeling still matters. Especially with AI.
- Power BI fundamentals still apply
- Not everything needs Fabric on day one

What to do next
- Understand where Fabric fits in your context and skills
- Explore features that make sense for you (trial capacity)



Thank you!

Any questions?

Session Feedback



Captain Platform