#FABCONSQLCON2026

FABCON
Microsoft Fabric
COMMUNITY CONFERENCE

Chaos to Clarity

Governance and Adoption Lessons from Rolling Out Power BI and Fabric

Chris Aleman

Data Analytics Manager

#FABCONSQLCON2026

FABCON
Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON
Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026

Why This Session

⬆️ Power BI + Fabric adoption scales fast

⬇️ Governance often lags behind adoption

🔗 Result: confusion, rework, loss of trust

Agenda

1. Setting the Stage: Chaos at Scale
   - What breaks down as Power BI and Fabric adoption grows
2. Governance Lessons Learned
   - Workspaces, permissions, and distribution patterns that didn’t scale
3. From Chaos to Clarity
   - Role-based access, apps, and enterprise semantic models
4. Driving Adoption with Guardrails
   - How governance accelerates trust and usage (instead of slowing it down)
5. AI-Ready by Design
   - Preparing data and semantic models for Copilot and future AI
6. Key Takeaways, Resources, & Q&A

What “Chaos” Looks Like At Scale

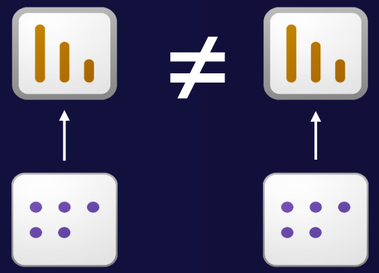

Multiple Versions of “truth”

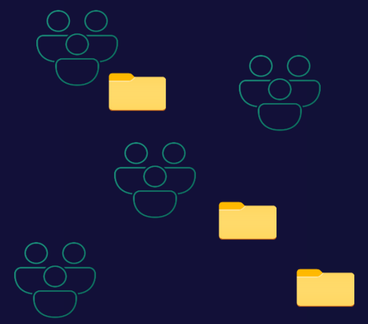

Fragmented Workspaces

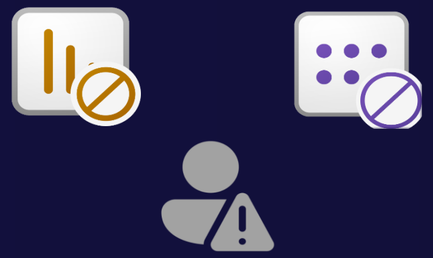

Messy Permissions

Frustrated end users

Governance Is About Behavior

- Governance ≠ locking everything down
- Governance = guiding how people work with data
- Balance empowerment and guardrails

Our Starting Point

Early Mistakes We Made

We moved fast. Then we paid the price

- Direct workspace access for consumers
- Department-level security groups
- No clear separation of build vs consume
- Consultants learning alongside us

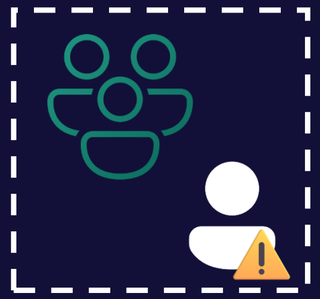

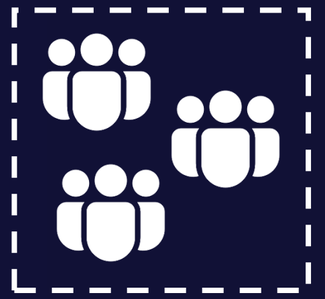

Why Workspaces Matter

Overall Architecture: How We Put It Together

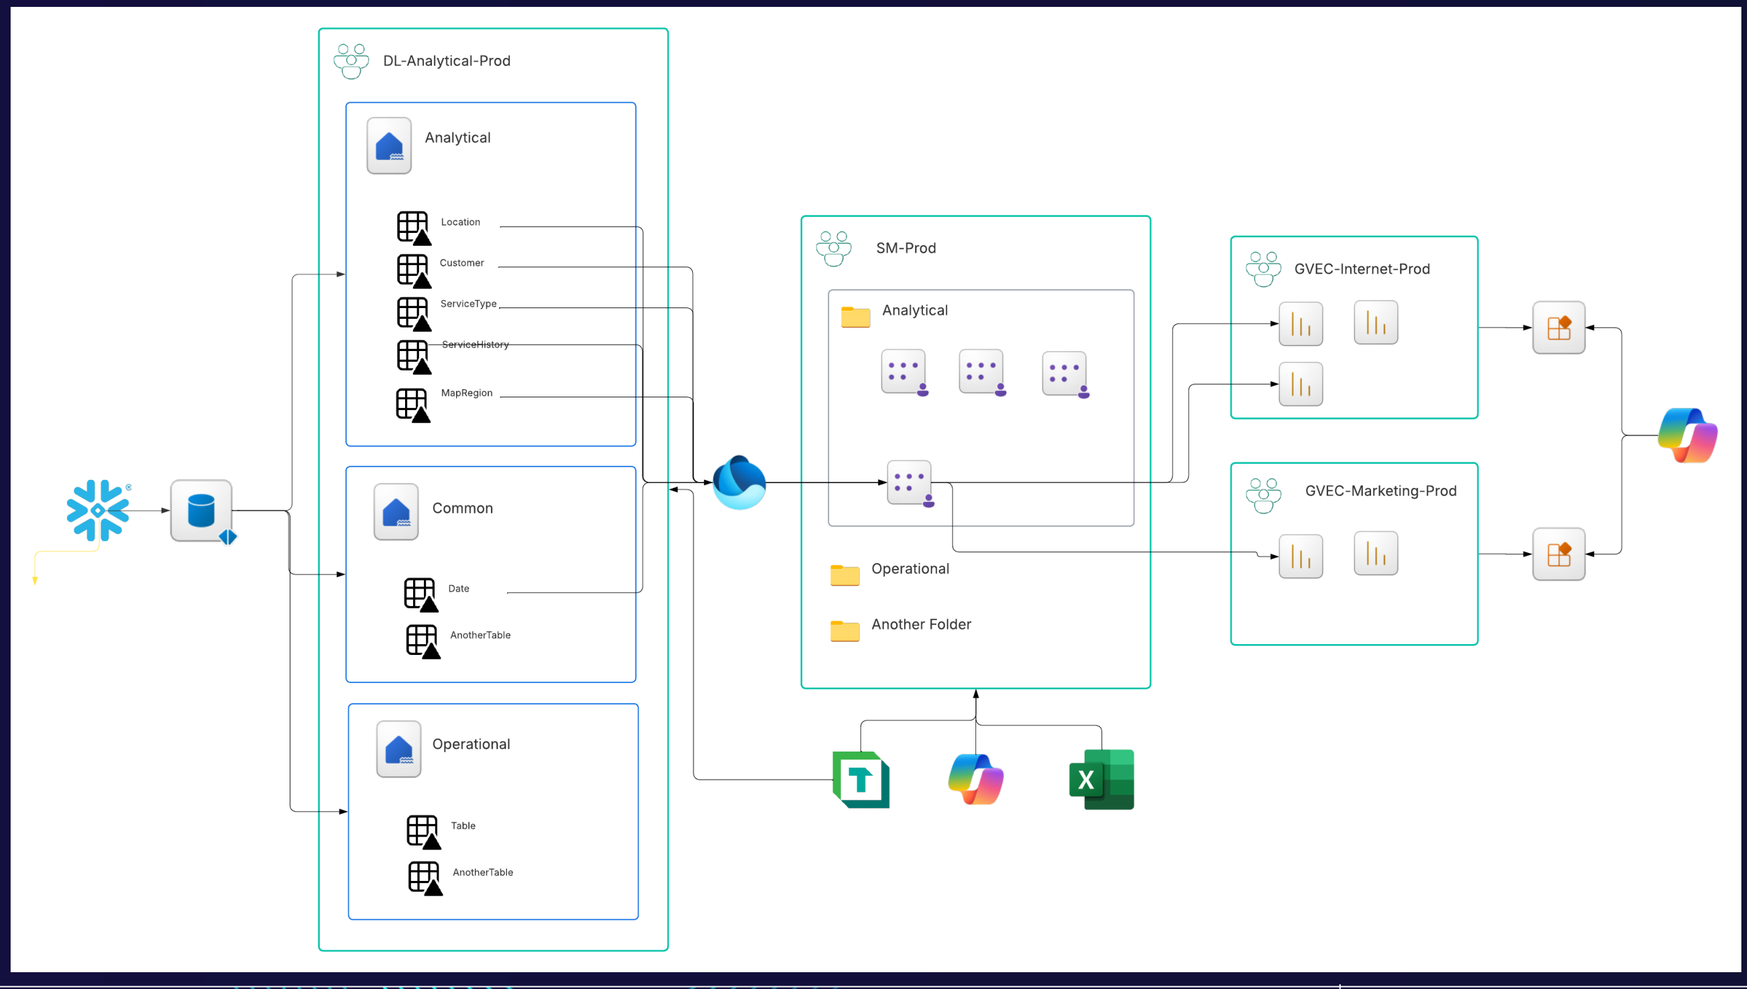

Workspace Strategy We Landed On

- Workspaces = build layer only
- Creators collaborate here
- Consumers do not live in workspaces

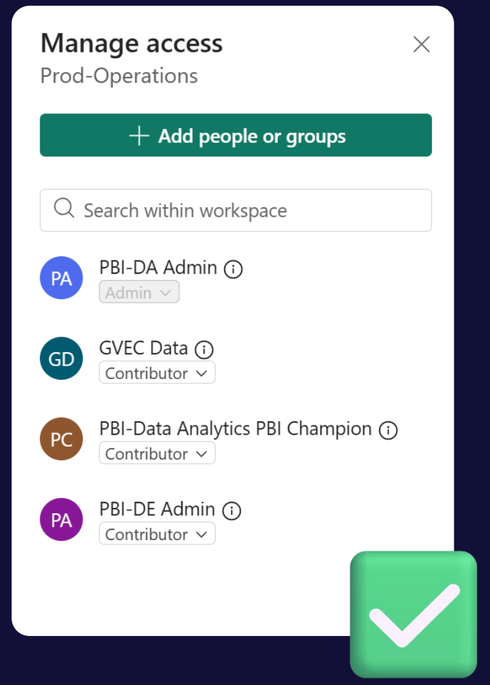

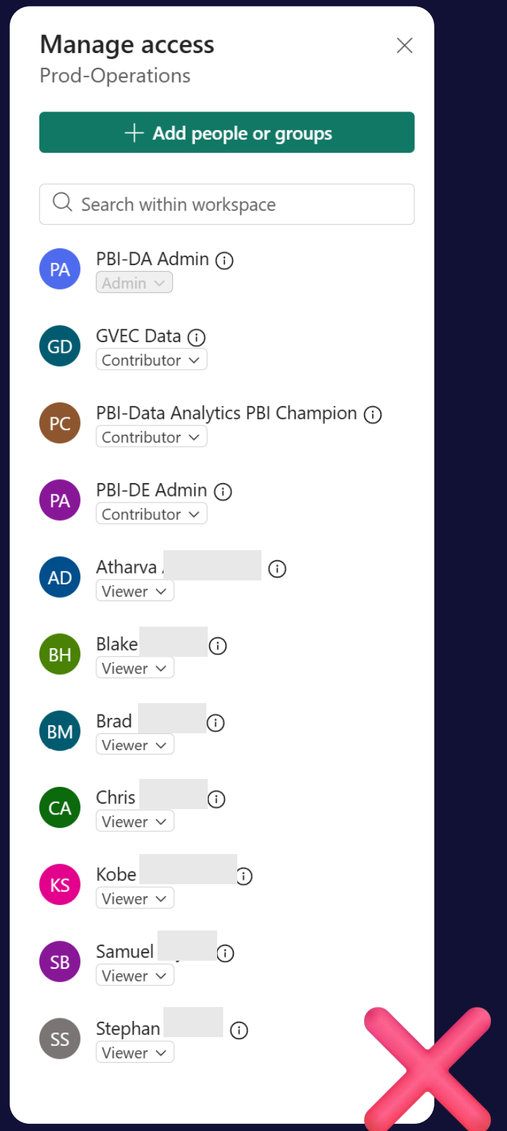

Apps as the Consumption Layer

- Apps = trusted entry point
- Clean navigation

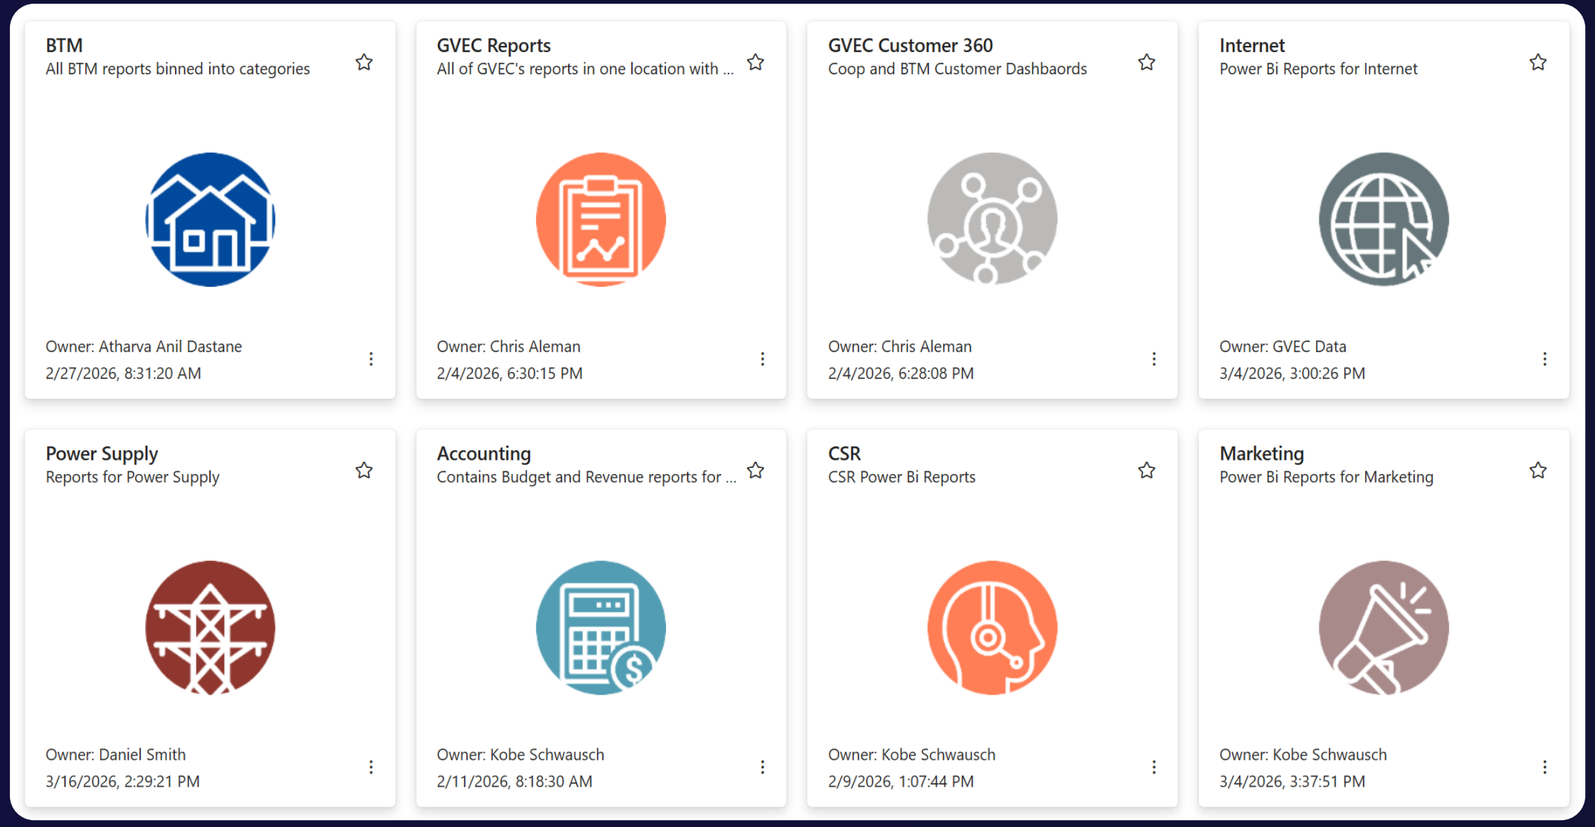

Apps as the Consumption Layer: Before

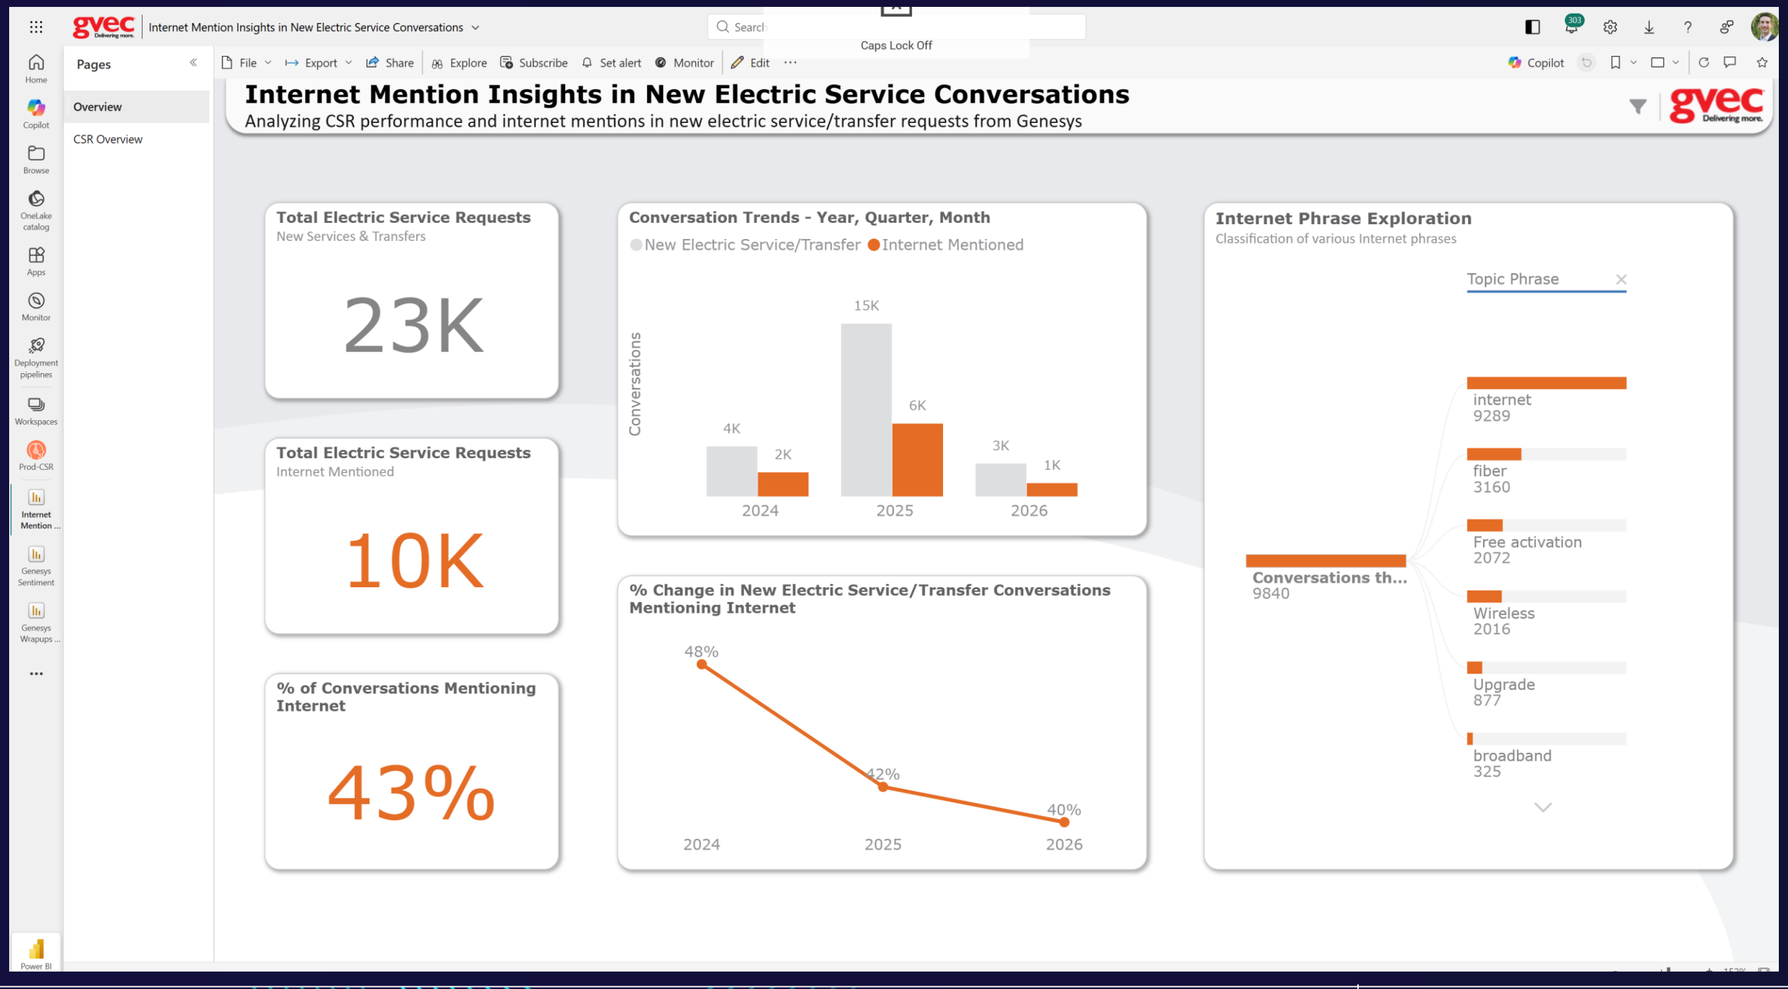

Apps as the Consumption Layer: After

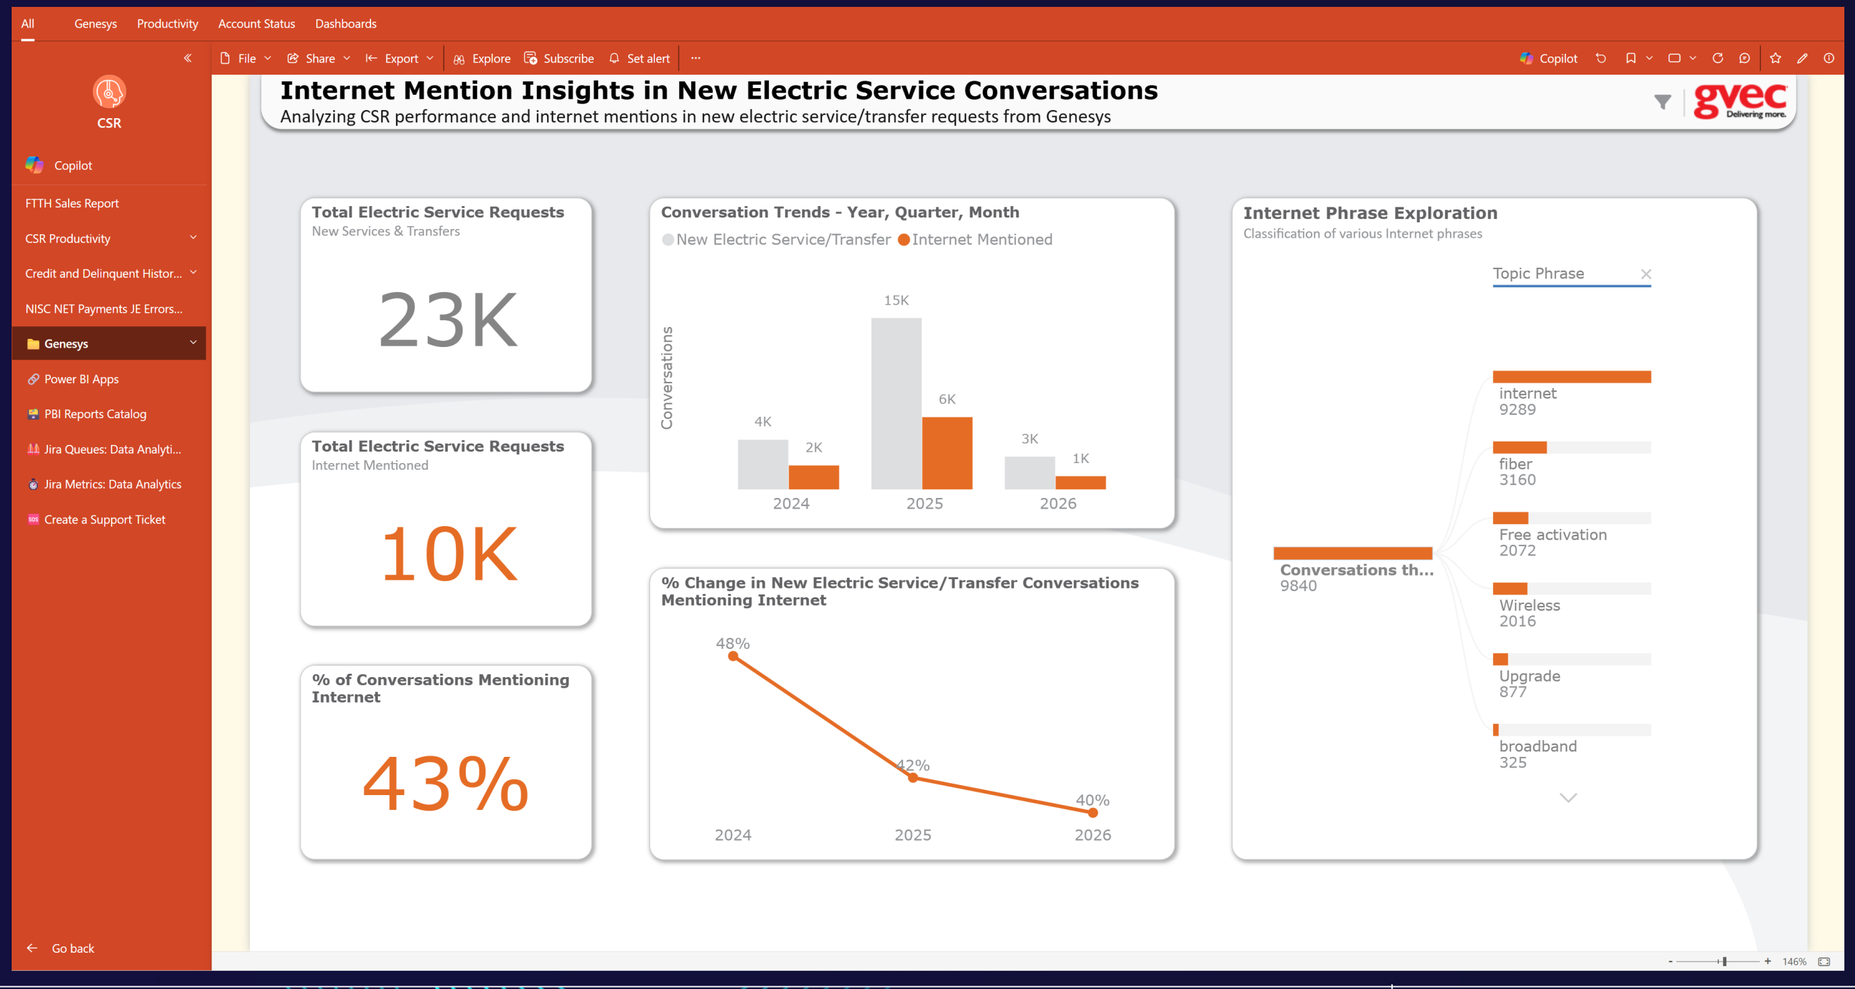

Apps as the Consumption Layer

- Role-based audiences
- Scales adoption cleanly

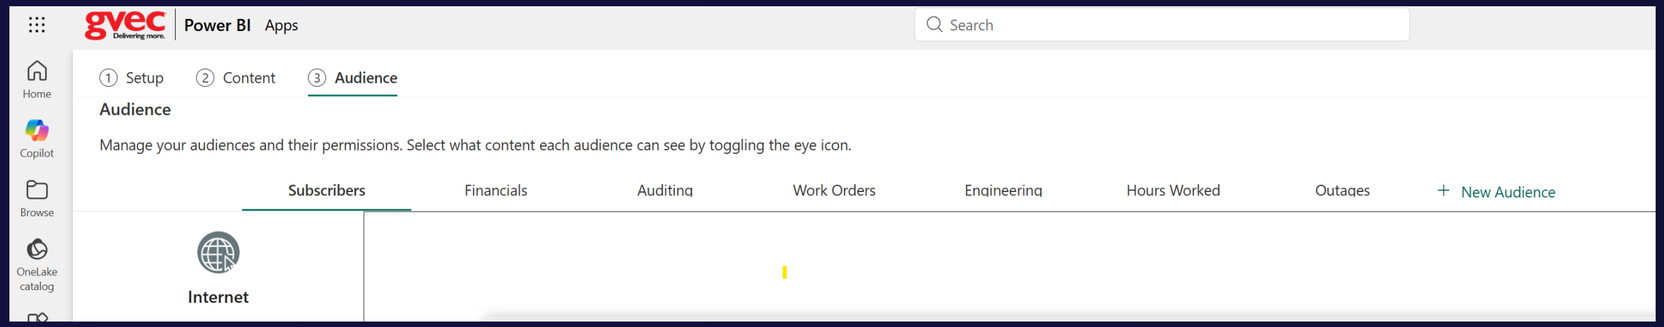

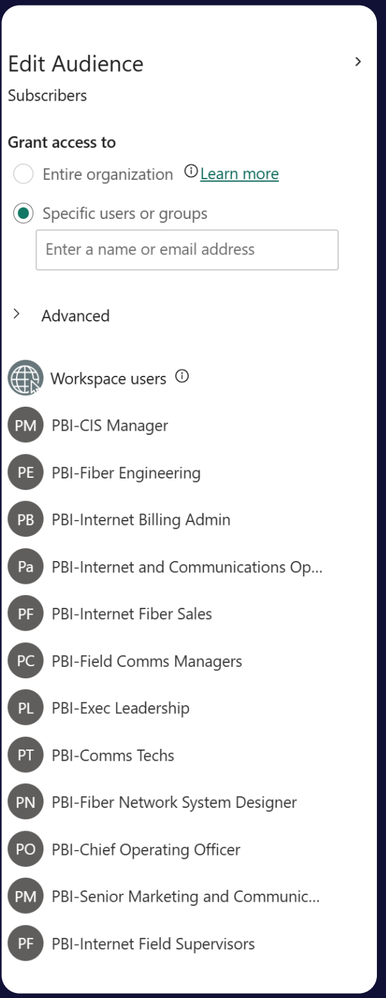

Role-Based Access (What Actually Works)

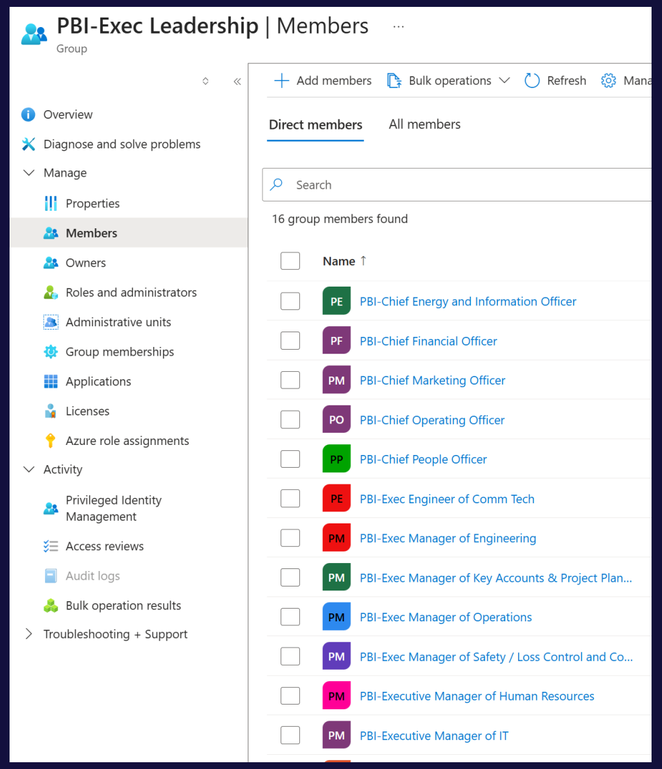

Entra ID security groups

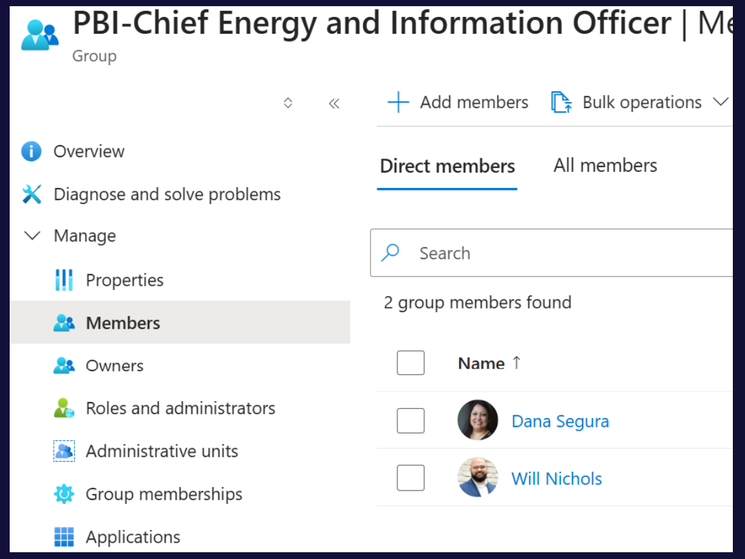

Group by job role, not department

Role-Based Access (What Actually Works)

- Dynamic membership where possible
- Fewer exceptions over time

And/Or	Property	Operator	Value
	jobTitle	Equals	Chief Energy and Information Officer
Or	jobTitle	Equals	Executive Assistant to C.E.I.O.

+ Add expression + Get custom extension properties ⓘ

The Semantic Model Shift

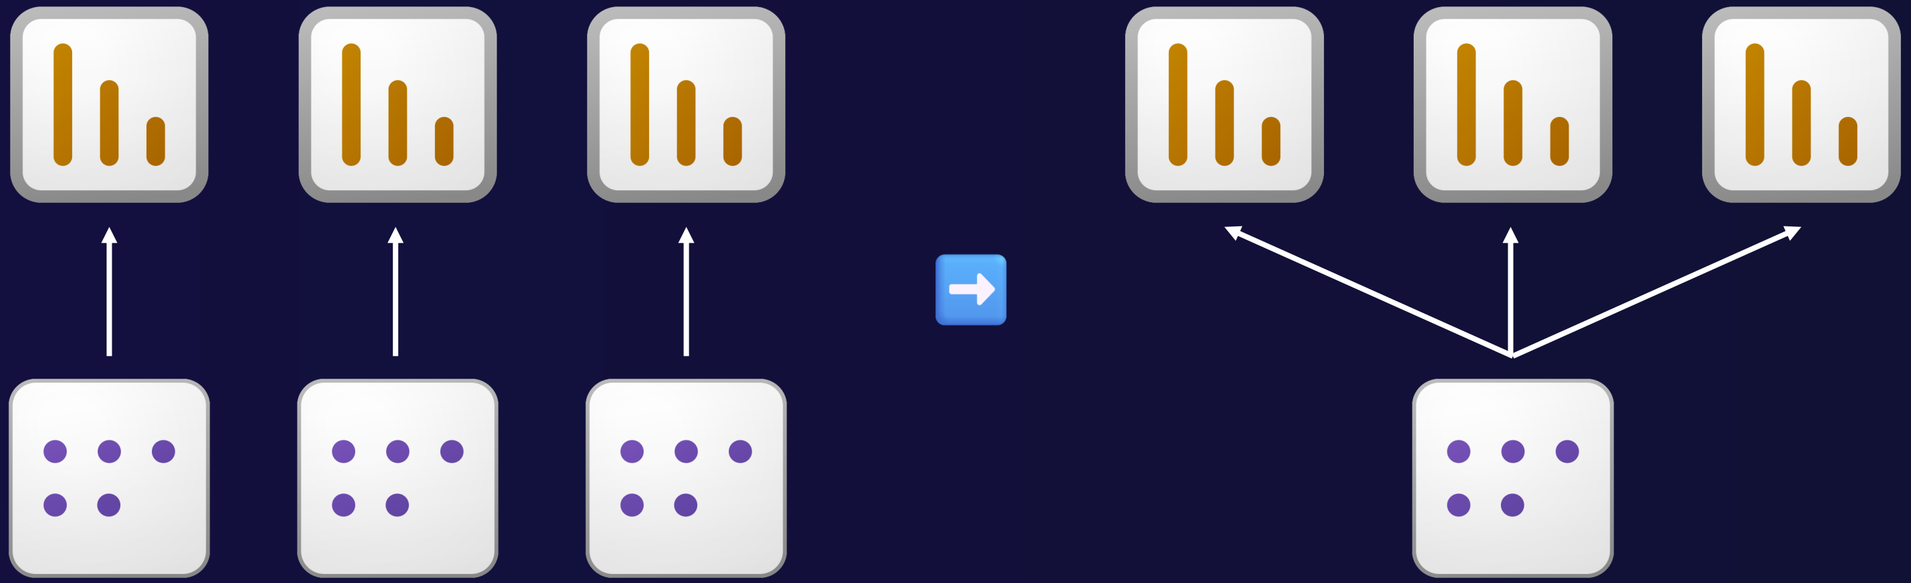

Semantic Model - Before

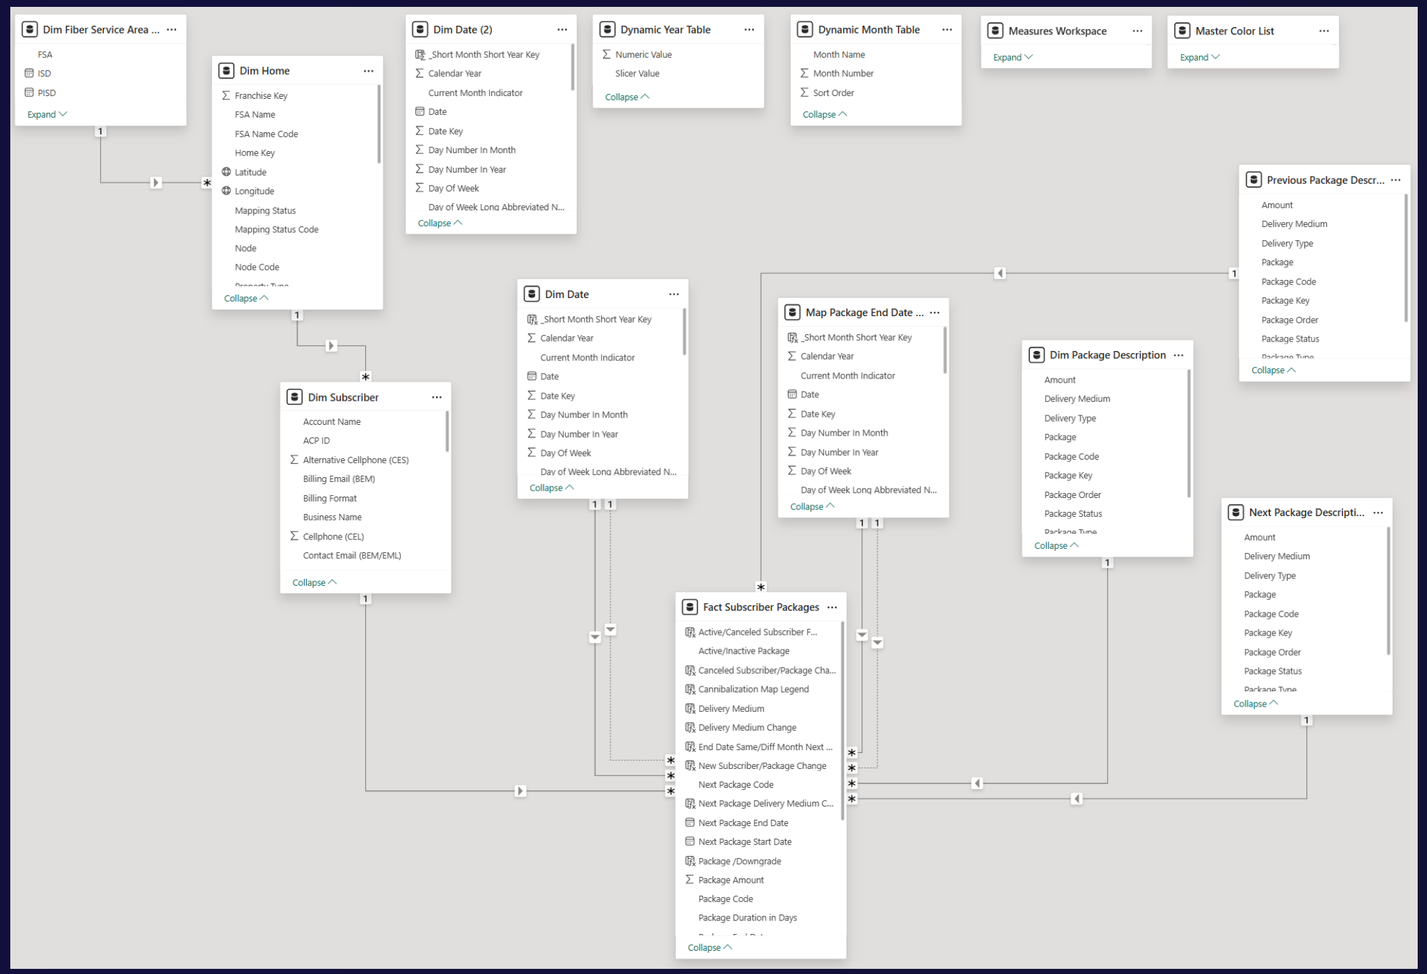

Semantic Model - After

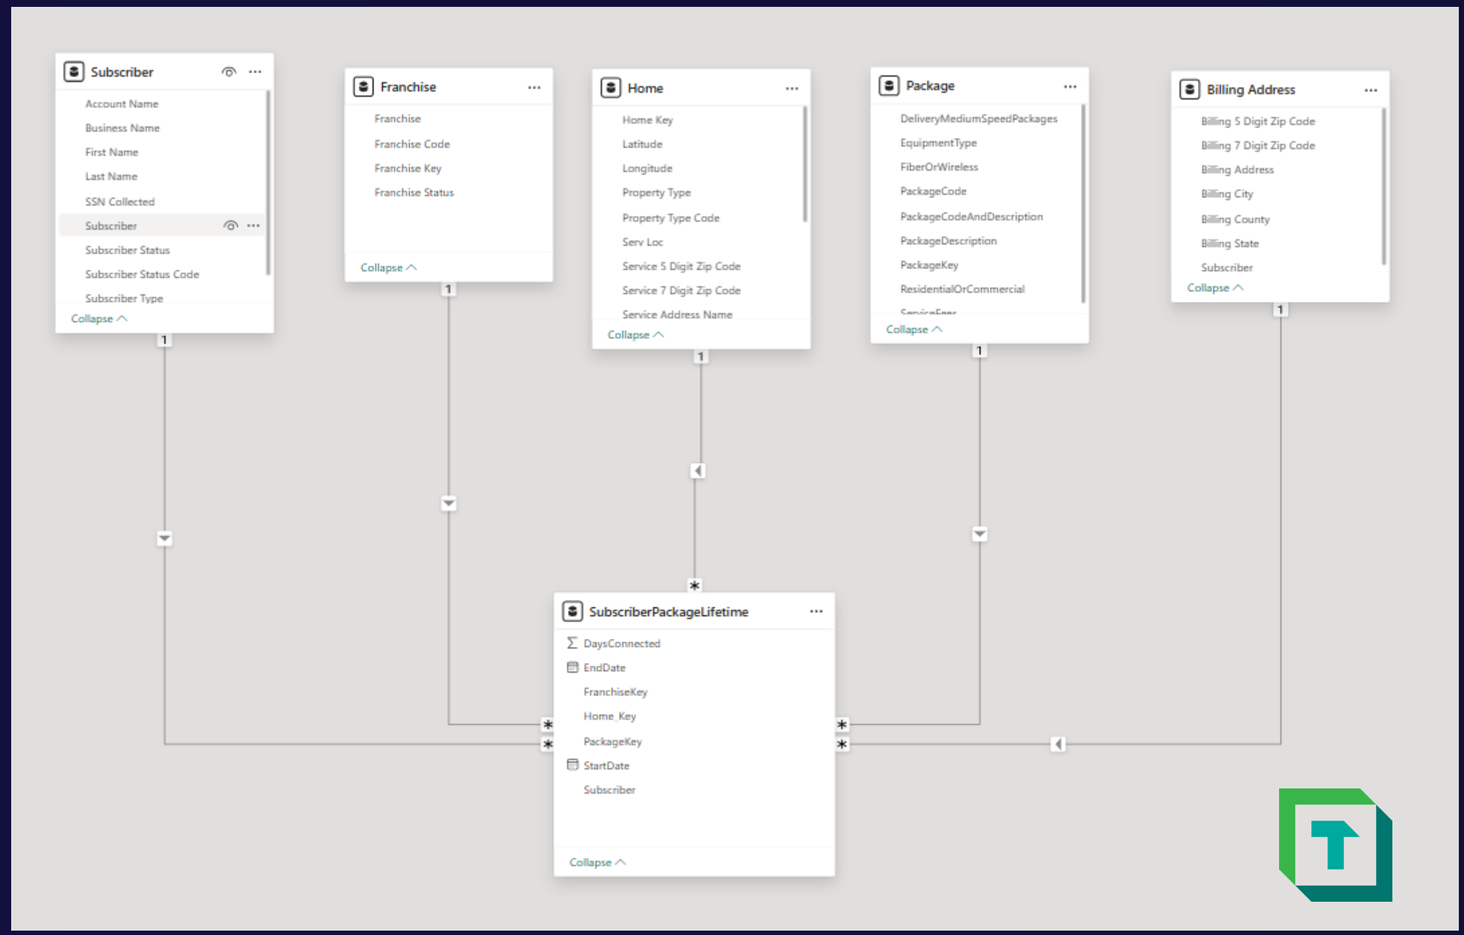

Enterprise Semantic Models

- Single source of truth tables
- Shared semantic models
- Thin reports
- Excel + Power BI + NLQ supported

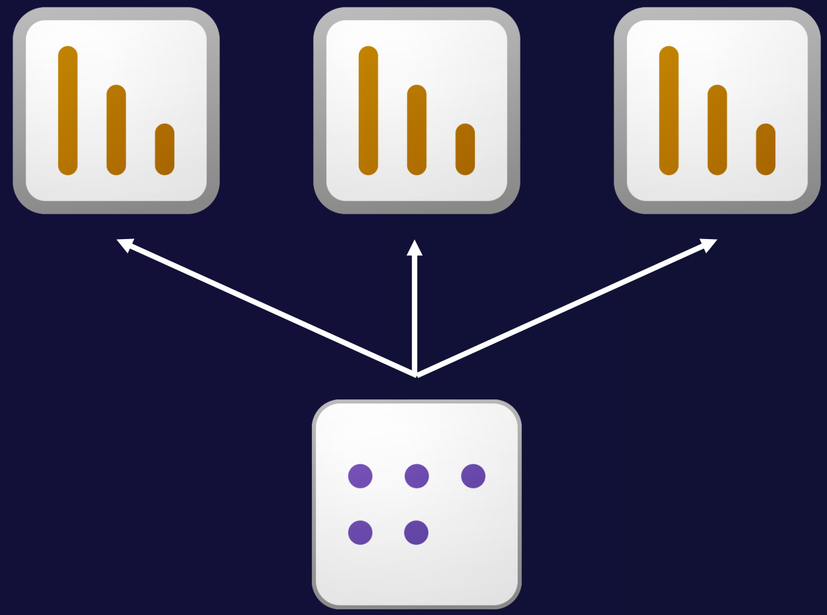

Source Control & Deployment

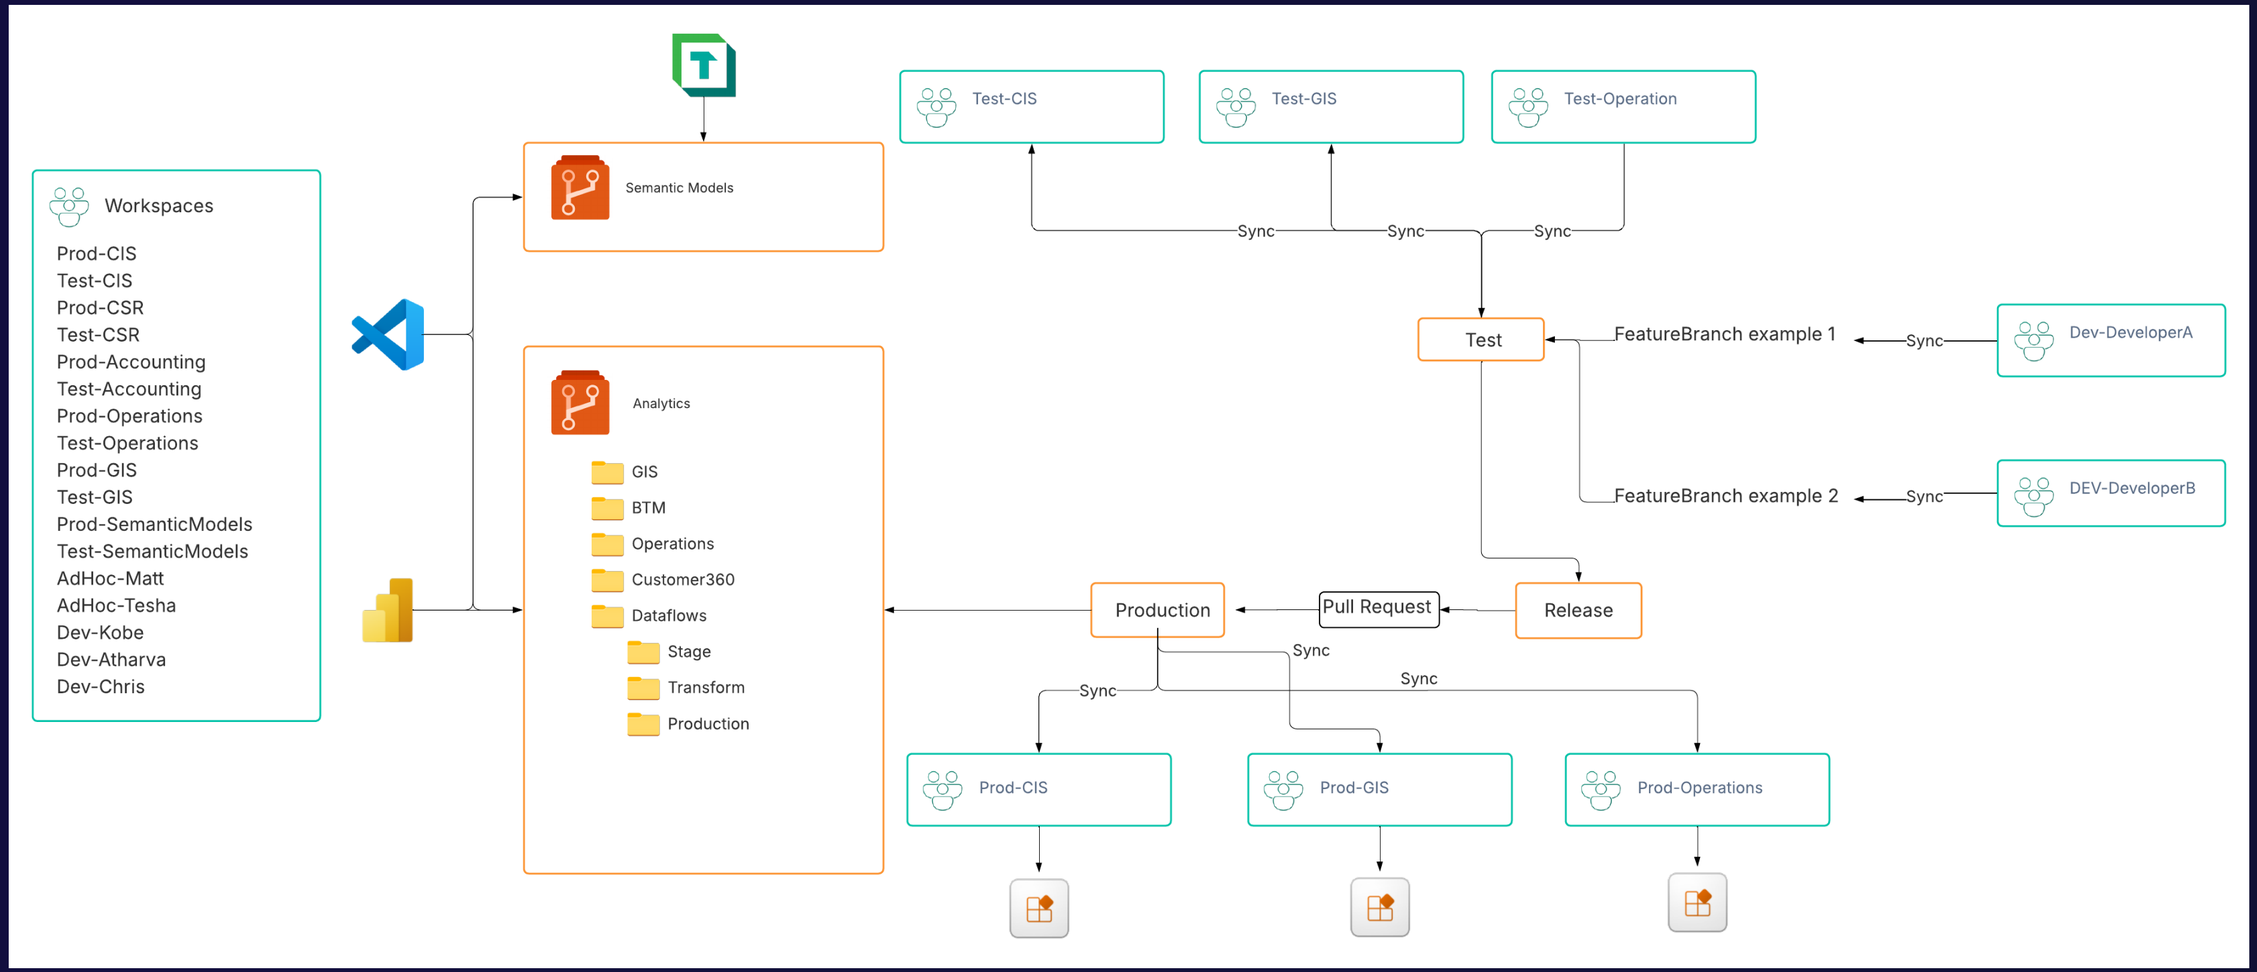

Governing Semantic Models

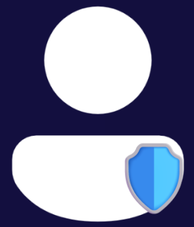

Ownership is explicit

Changes are intentional

Quality is enforced
before users feel pain

Dataflows: Bridge, Not Destination

- Dataflows Gen2 helped early
- Some business logic still lives there
- Long-term goal: push logic upstream
- Reduce report-side transformations

Monitoring & Visibility

🛡️ Oversharing is a governance risk

👁️ Need visibility into:

👥 Sharing patterns

🗂️ Workspace sprawl

👤➖ Orphaned assets

Monitoring & Visibility

- Adoption signals (views, unique users)
- Refresh failures / reliability
- Semantic model duplication
- Workspace sprawl
- Sharing patterns (direct share, link sprawl)

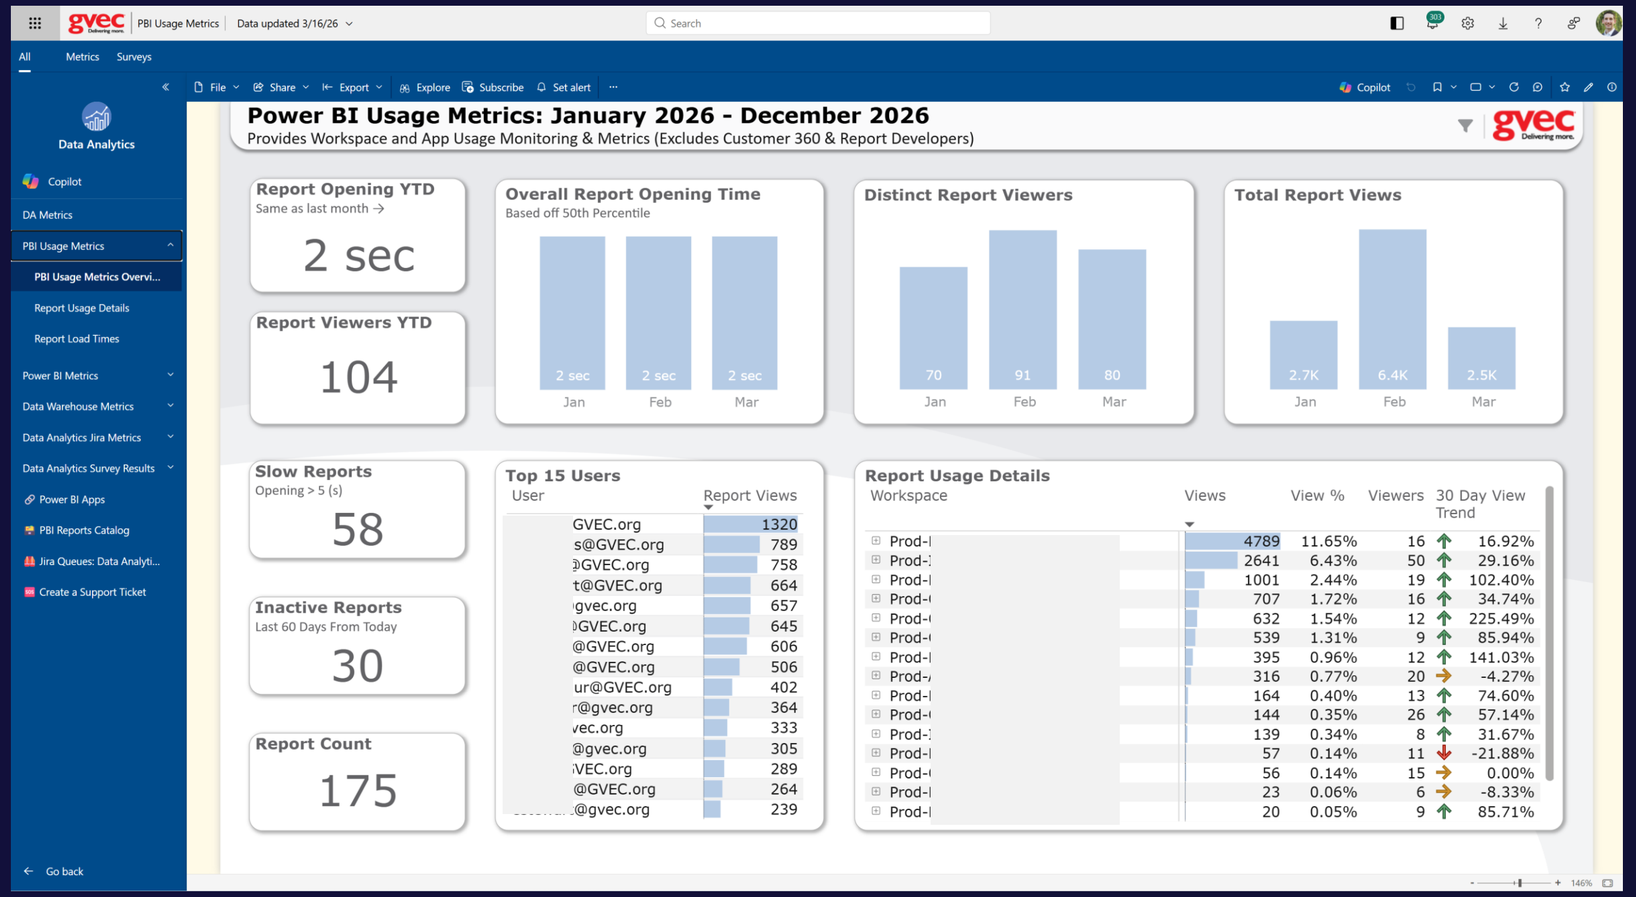

Monitoring & Visibility

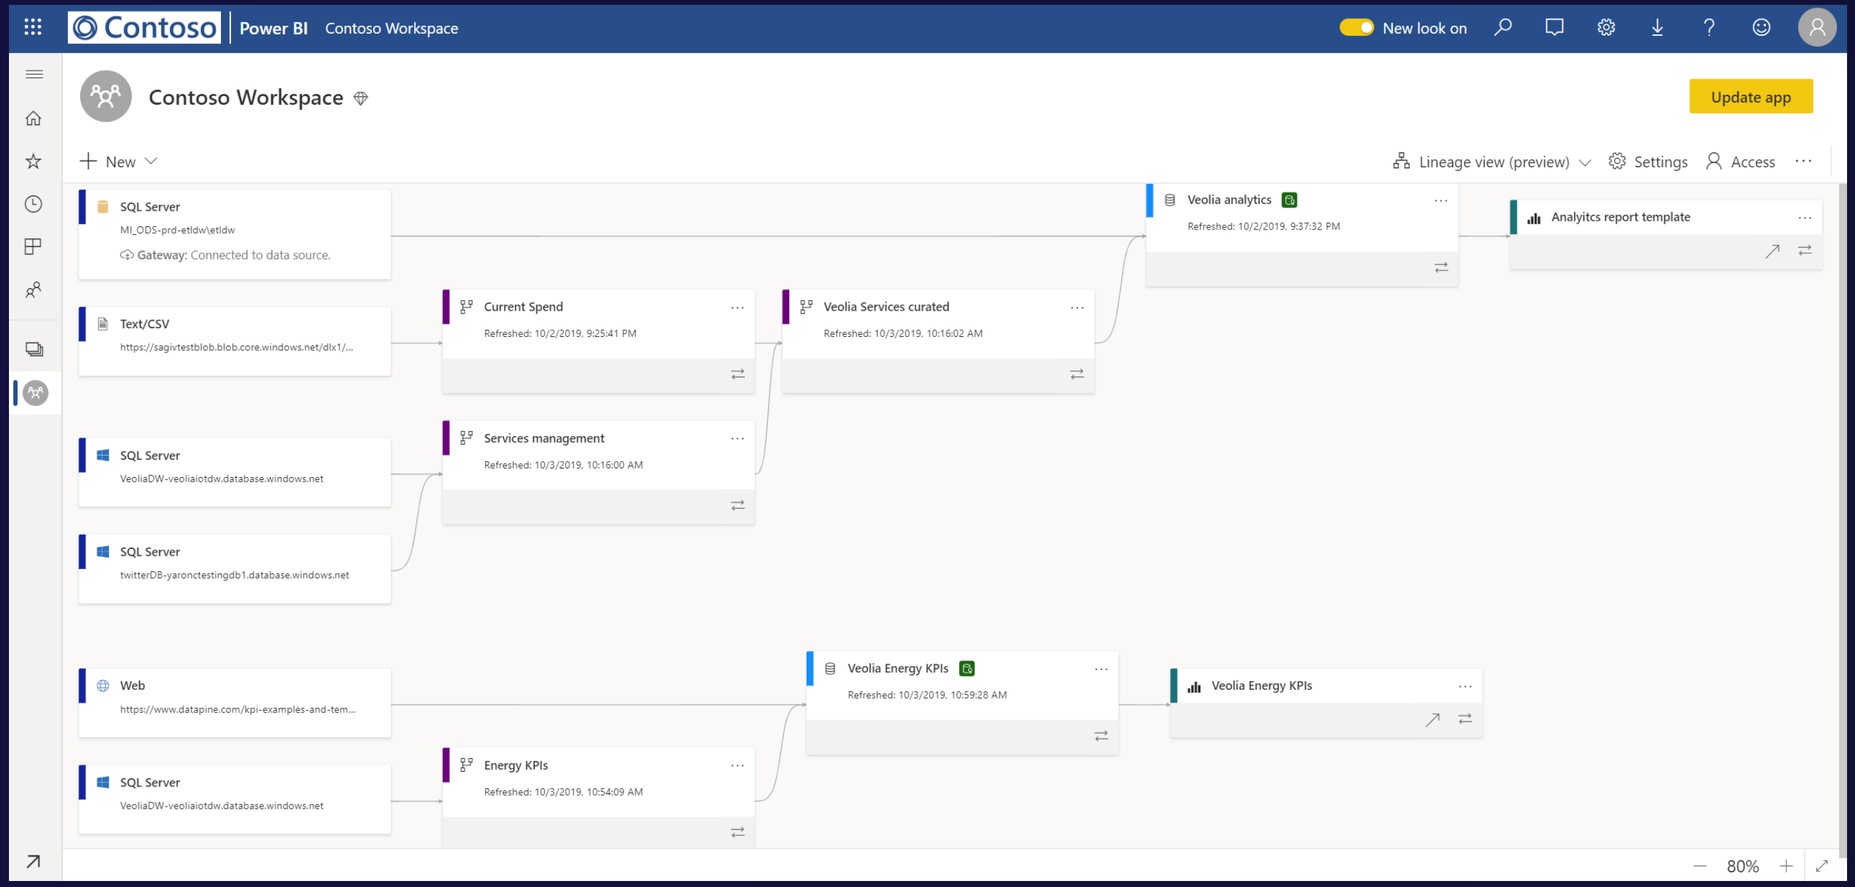

Adoption ≠ Enablement

- Users need clarity, not just access
- Training and structure matter
- Apps reduced “where do I go?” questions

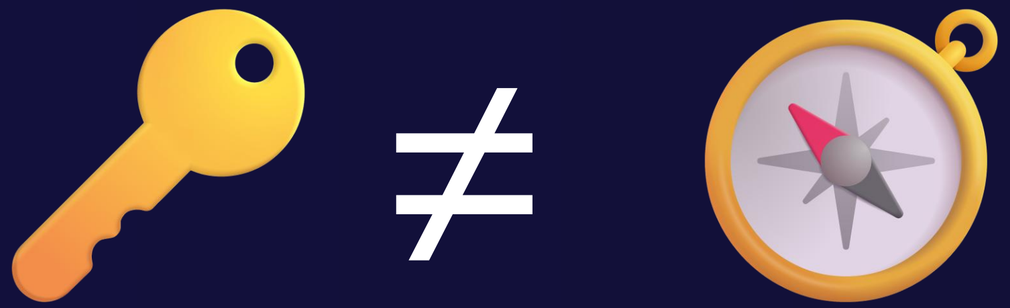

Access ≠ Enablement

Why AI Raises the Stakes

🗄️ Data Quality

🤖 AI / Copilot

✅ / ⚠️ Answers

What “AI-Ready” Actually Means

AI / Copilot

Definitions

Grain

Measures

Access

If We Did It Again

- ❄️ Go straight to Snowflake sooner
- 🛡️ Design security groups up front
- 🧭 Apps from day one
- 🔁 Fewer like-for-like conversions

Trust Is the Real KPI

- Trust is slow to build
- Easy to lose
- Consistency beats speed long-term
- Identify business champions early

Outcomes

Consistency

Governance

Trust

Governance Playbook (Summary)

- 👁️ Monitor Continuously
- ✅ Enforce Quality Early
- 🧠 Shared Semantic Models
- 🔐 Role-Based Security Groups
- 🧭 Apps as Default Distribution
- 🧱 Separate Build / Consume

Key Takeaways

⚠️ Chaos is predictable without structure

🛡️ Governance enables adoption

📄 Semantic models are the contract

🧭 AI rewards discipline, not shortcuts

Trusted Resources

- Microsoft Learn Fabric Adoption Roadmap
- Microsoft Learn Power Bi Implementation Planning
- sqlbi
- Tabular Editor
- Data Goblins
- Radacad
- Story Telling With Data
- Guy in a Cube
- Chris Webb’s Blog
- Powerbi.tips
- Paul Turley

Thank You!

Questions?

Chris Aleman

Caleman@gvec.org

Sound off.
The mic is all yours.
Influence the product roadmap.

Join the Fabric User Panel

Share your feedback directly with our Fabric product group and researchers.

https://aka.ms/JoinFabricUserPanel

Join the SQL User Panel

Influence our SQL roadmap and ensure it meets your real-life needs

https://aka.ms/JoinSQLUserPanel

Get Two Fabric Certifications for FREE

Attendees of FABCON can take the Fabric Analytics Engineer or Fabric Data Engineer exam for free. Be part of the 2 fastest growing role-based certifications in Microsoft history.

Request your voucher by March 23, 2026.

https://aka.ms/fabcon/cert100

How was the session?

Complete Session Surveys in Whova for your chance to WIN PRIZES!

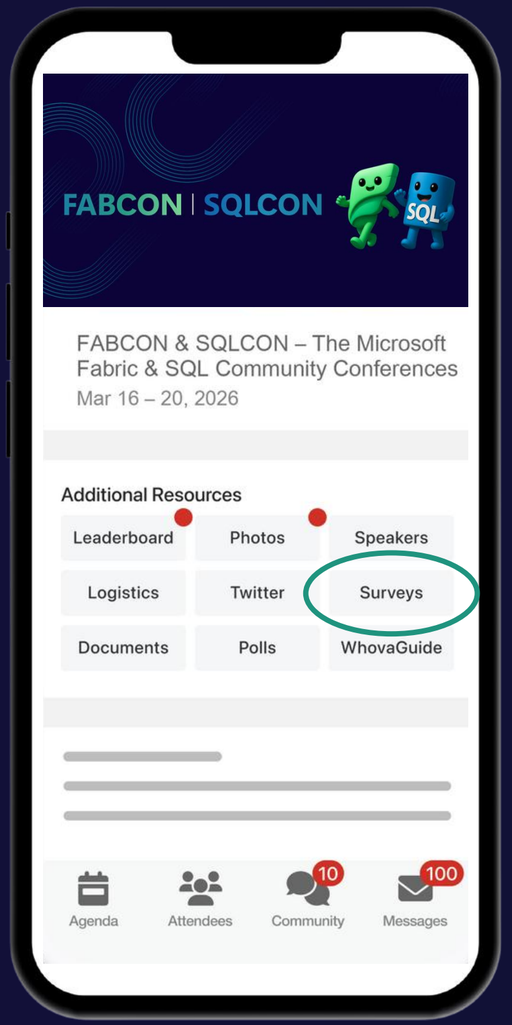

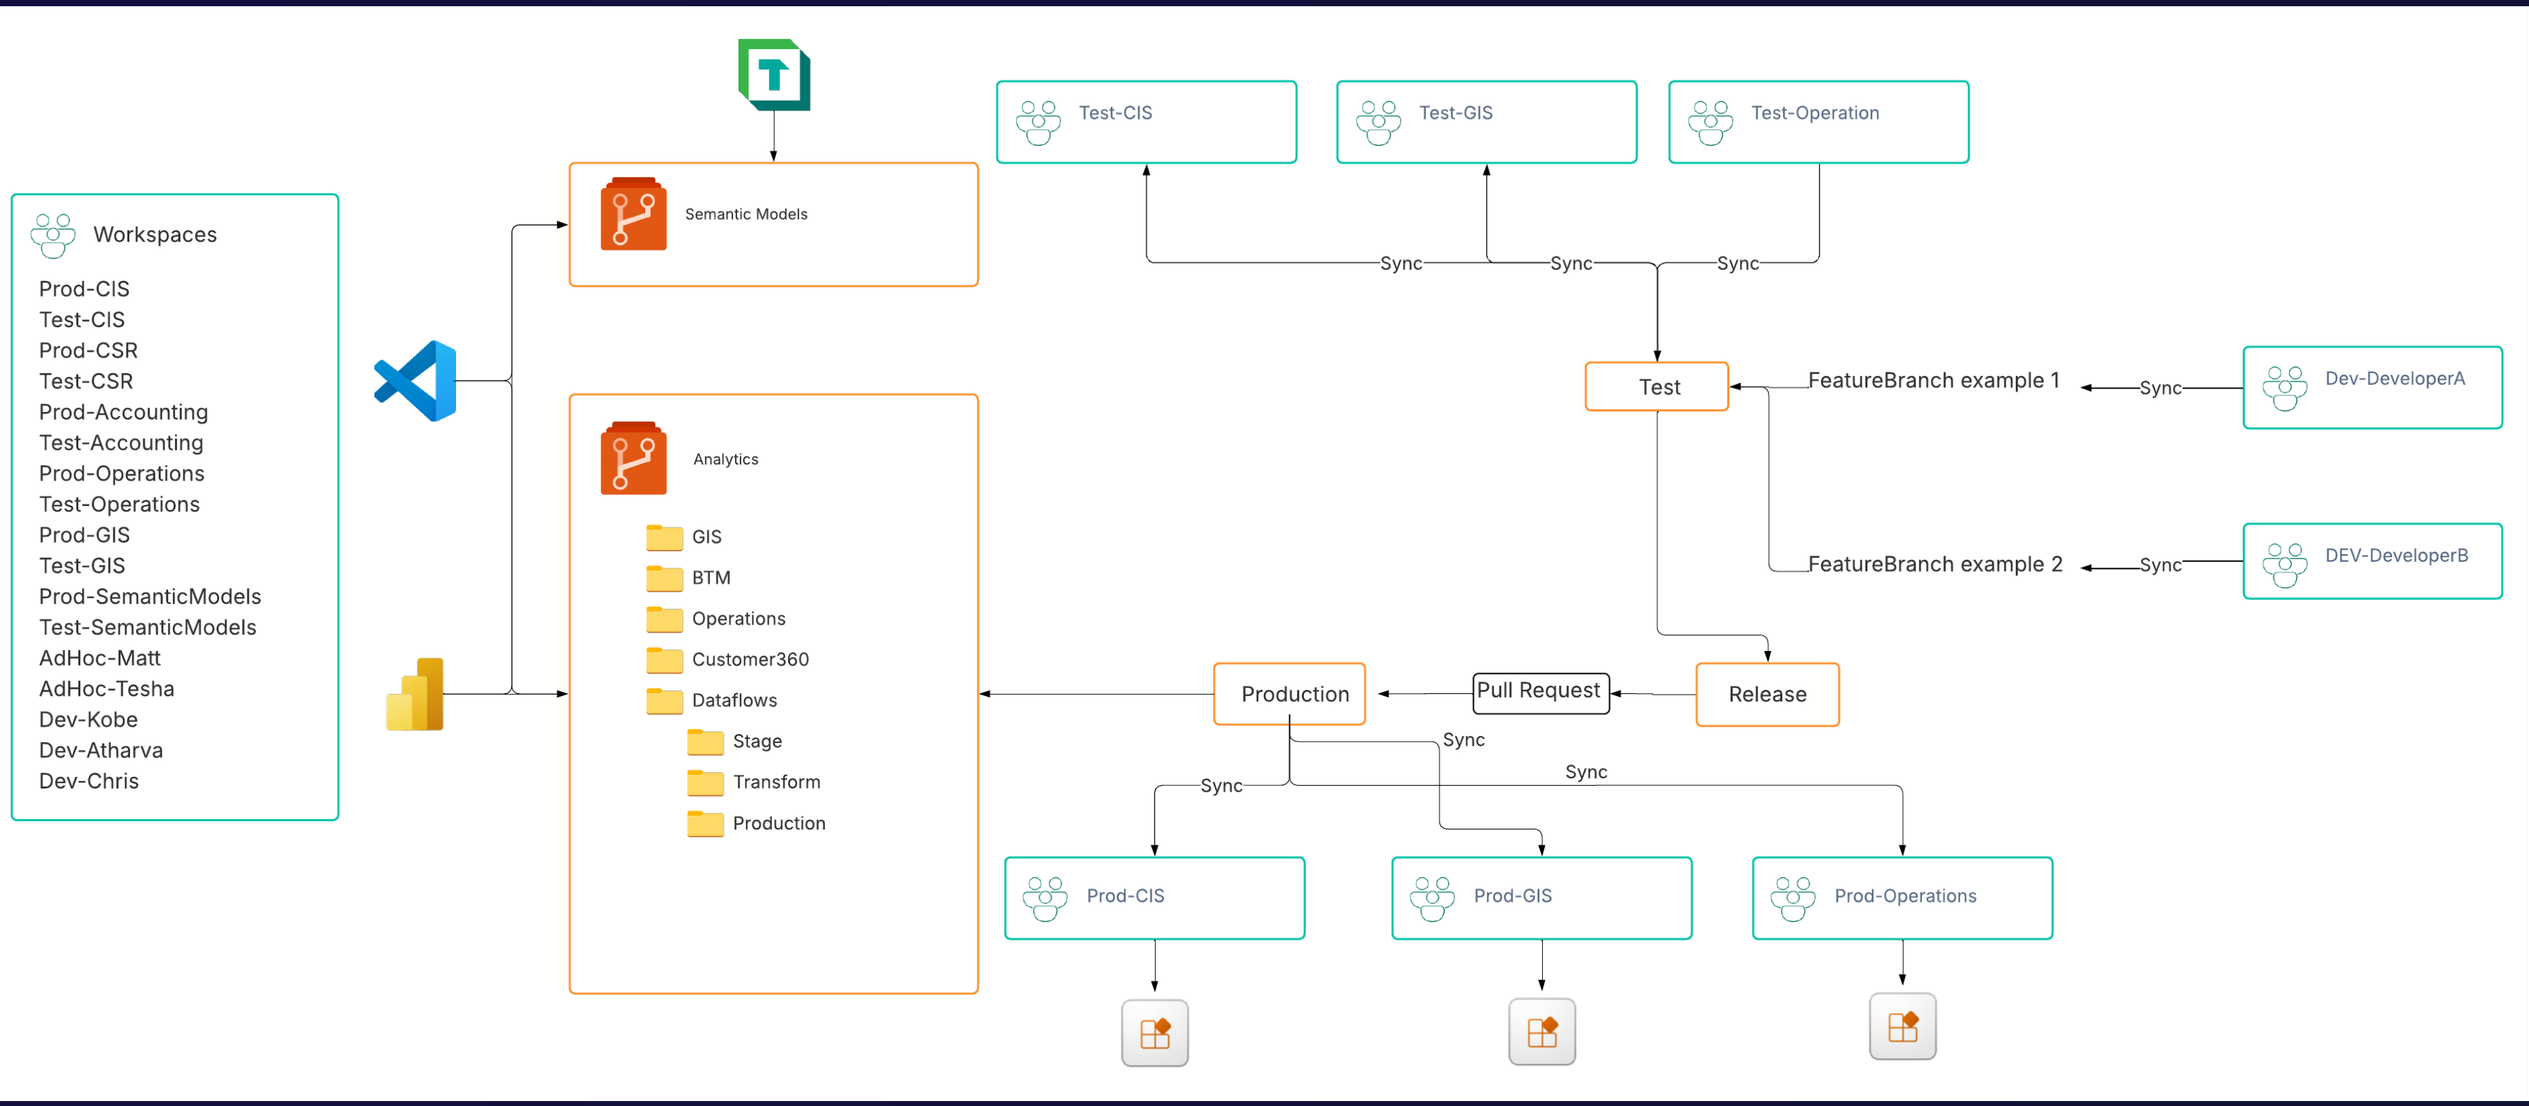
Workspaces
Prod-CIS
Test-CIS
Prod-CSR
Test-CSR
Prod-Accounting
Test-Accounting
Prod-Operations
Test-Operations
Prod-GIS
Test-GIS
Prod-SemanticModels
Test-SemanticModels
AdHoc-Matt
AdHoc-Tesha
Dev-Kobe
Dev-Atharva
Dev-Chris
Semantic Models
Analytics
GIS
BTM
Operations
Customer360
Dataflows
Stage
Transform
Production
Test-CIS
Test-GIS
Test-Operation
Sync
Sync
Sync
Test
FeatureBranch example 1
Sync
Dev-DeveloperA
FeatureBranch example 2
Sync
DEV-DeveloperB
Release
Pull Request
Production
Sync
Sync
Sync
Prod-CIS
Prod-GIS
Prod-Operations

Semantic Model Repository
- All semantic models will be stored in a dedicated repo.
- Analysts and engineers will have access to production and test branches.
- Once a certification process is defined, Prod branch access will be restricted.

Tabular Editor will be the primary tool for creating and maintaining semantic models. It is a Microsoft-approved external tool and is included in the Fabric Engineering and Fabric Solution Architecture certifications.

Test Workspaces
- Each test workspace in Fabric will be synced with it's corresponding folder in the Analytics repo
- When end users need to test a report, they will be given direct access to teh repor in the appropriate test workspace. After a change has been pushed to prod, permissions will need to be revoked.
  - Periodic access auditing will be needed.
- When changes are made in the repo, the Fabric workspace will need to resync to the repo- manual process at the moment but may be automated in the future.

Developer workspaces
- Available as an option if needed.
- Dev workspaces will be pointed to the current project a develope ris working in and is intended to be used only by the developer, not end users.

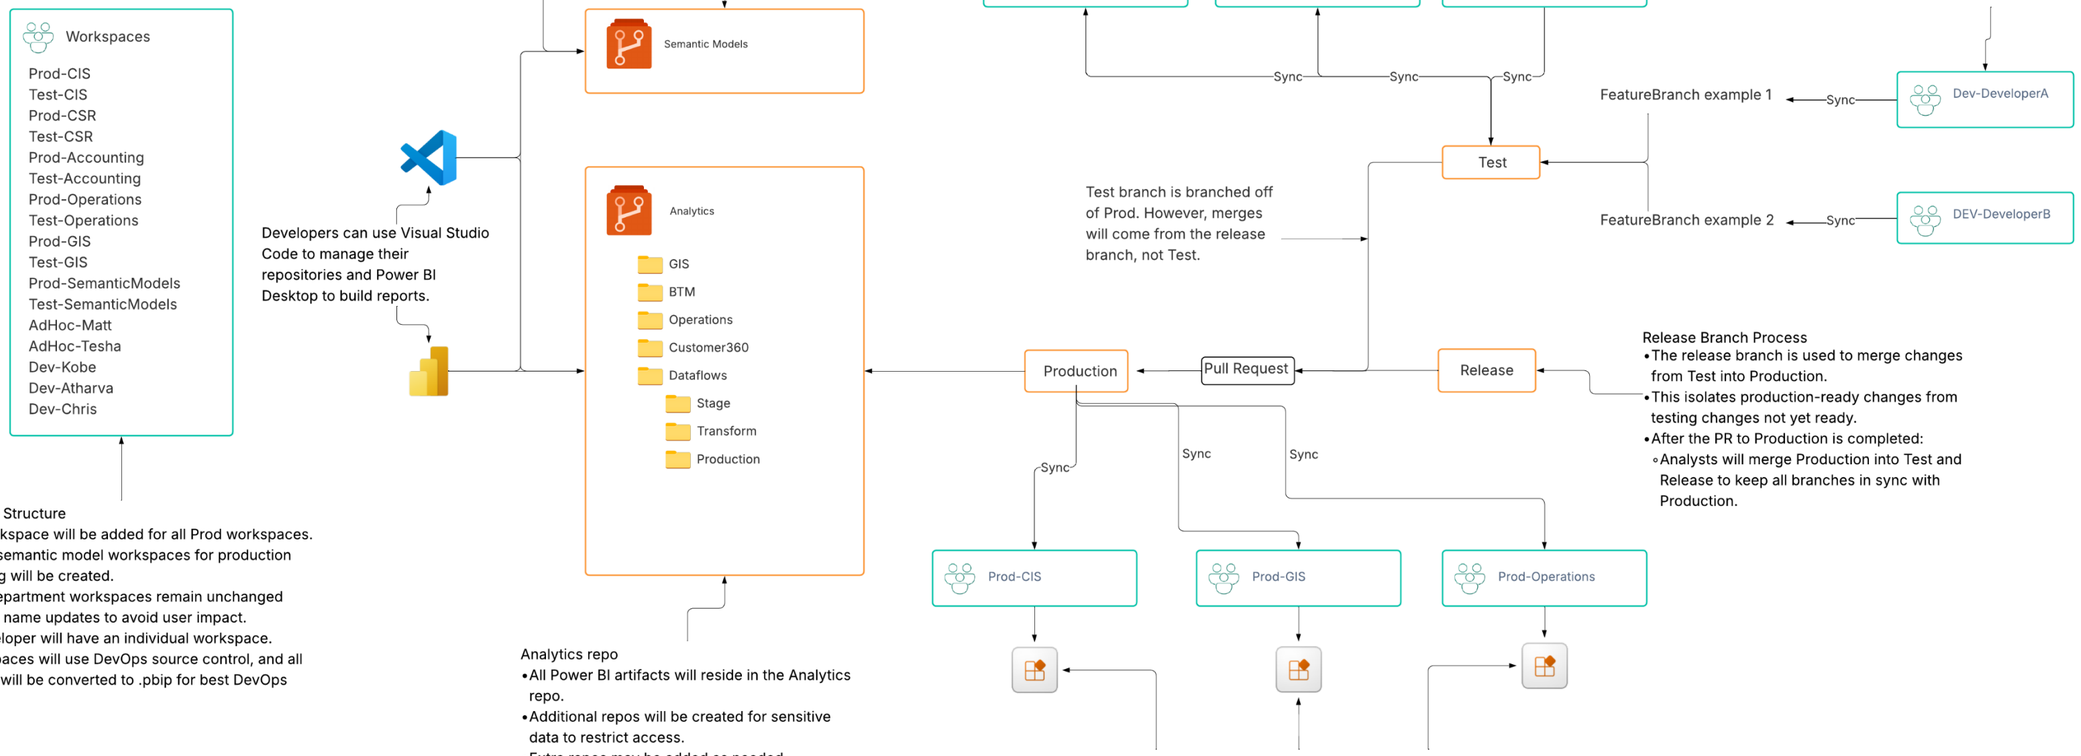

Developers can use Visual Studio Code to manage their repositories and Power BI Desktop to build reports.

Test branch is branched off of Prod. However, merges will come from the release branch, not Test.

Release Branch Process
- The release branch is used to merge changes from Test into Production.
- This isolates production-ready changes from testing changes not yet ready.
- After the PR to Production is completed:
  - Analysts will merge Production into Test and Release to keep all branches in sync with Production.

Workspace Structure
- A test workspace will be added for all Prod workspaces.
- Separate semantic model workspaces for production and testing will be created.
- Current department workspaces remain unchanged except for name updates to avoid user impact.
- Each developer will have an individual workspace.
- All workspaces will use DevOps source control, and all .pbix files will be converted to .pbip for best DevOps practices.

Analytics repo
- All Power BI artifacts will reside in the Analytics repo.
- Additional repos will be created for sensitive data to restrict access.
- Extra repos may be added as needed.
- Analysts and engineers will have access to the repo
- Once a certification process is defined, Prod branch access may be restricted.

Production Workspaces & Apps
- Existing Prod workspaces and apps remain unchanged to avoid user impact.
- After changes are synced to Prod, the app must still be updated.
- A future process will be created to automate app updates

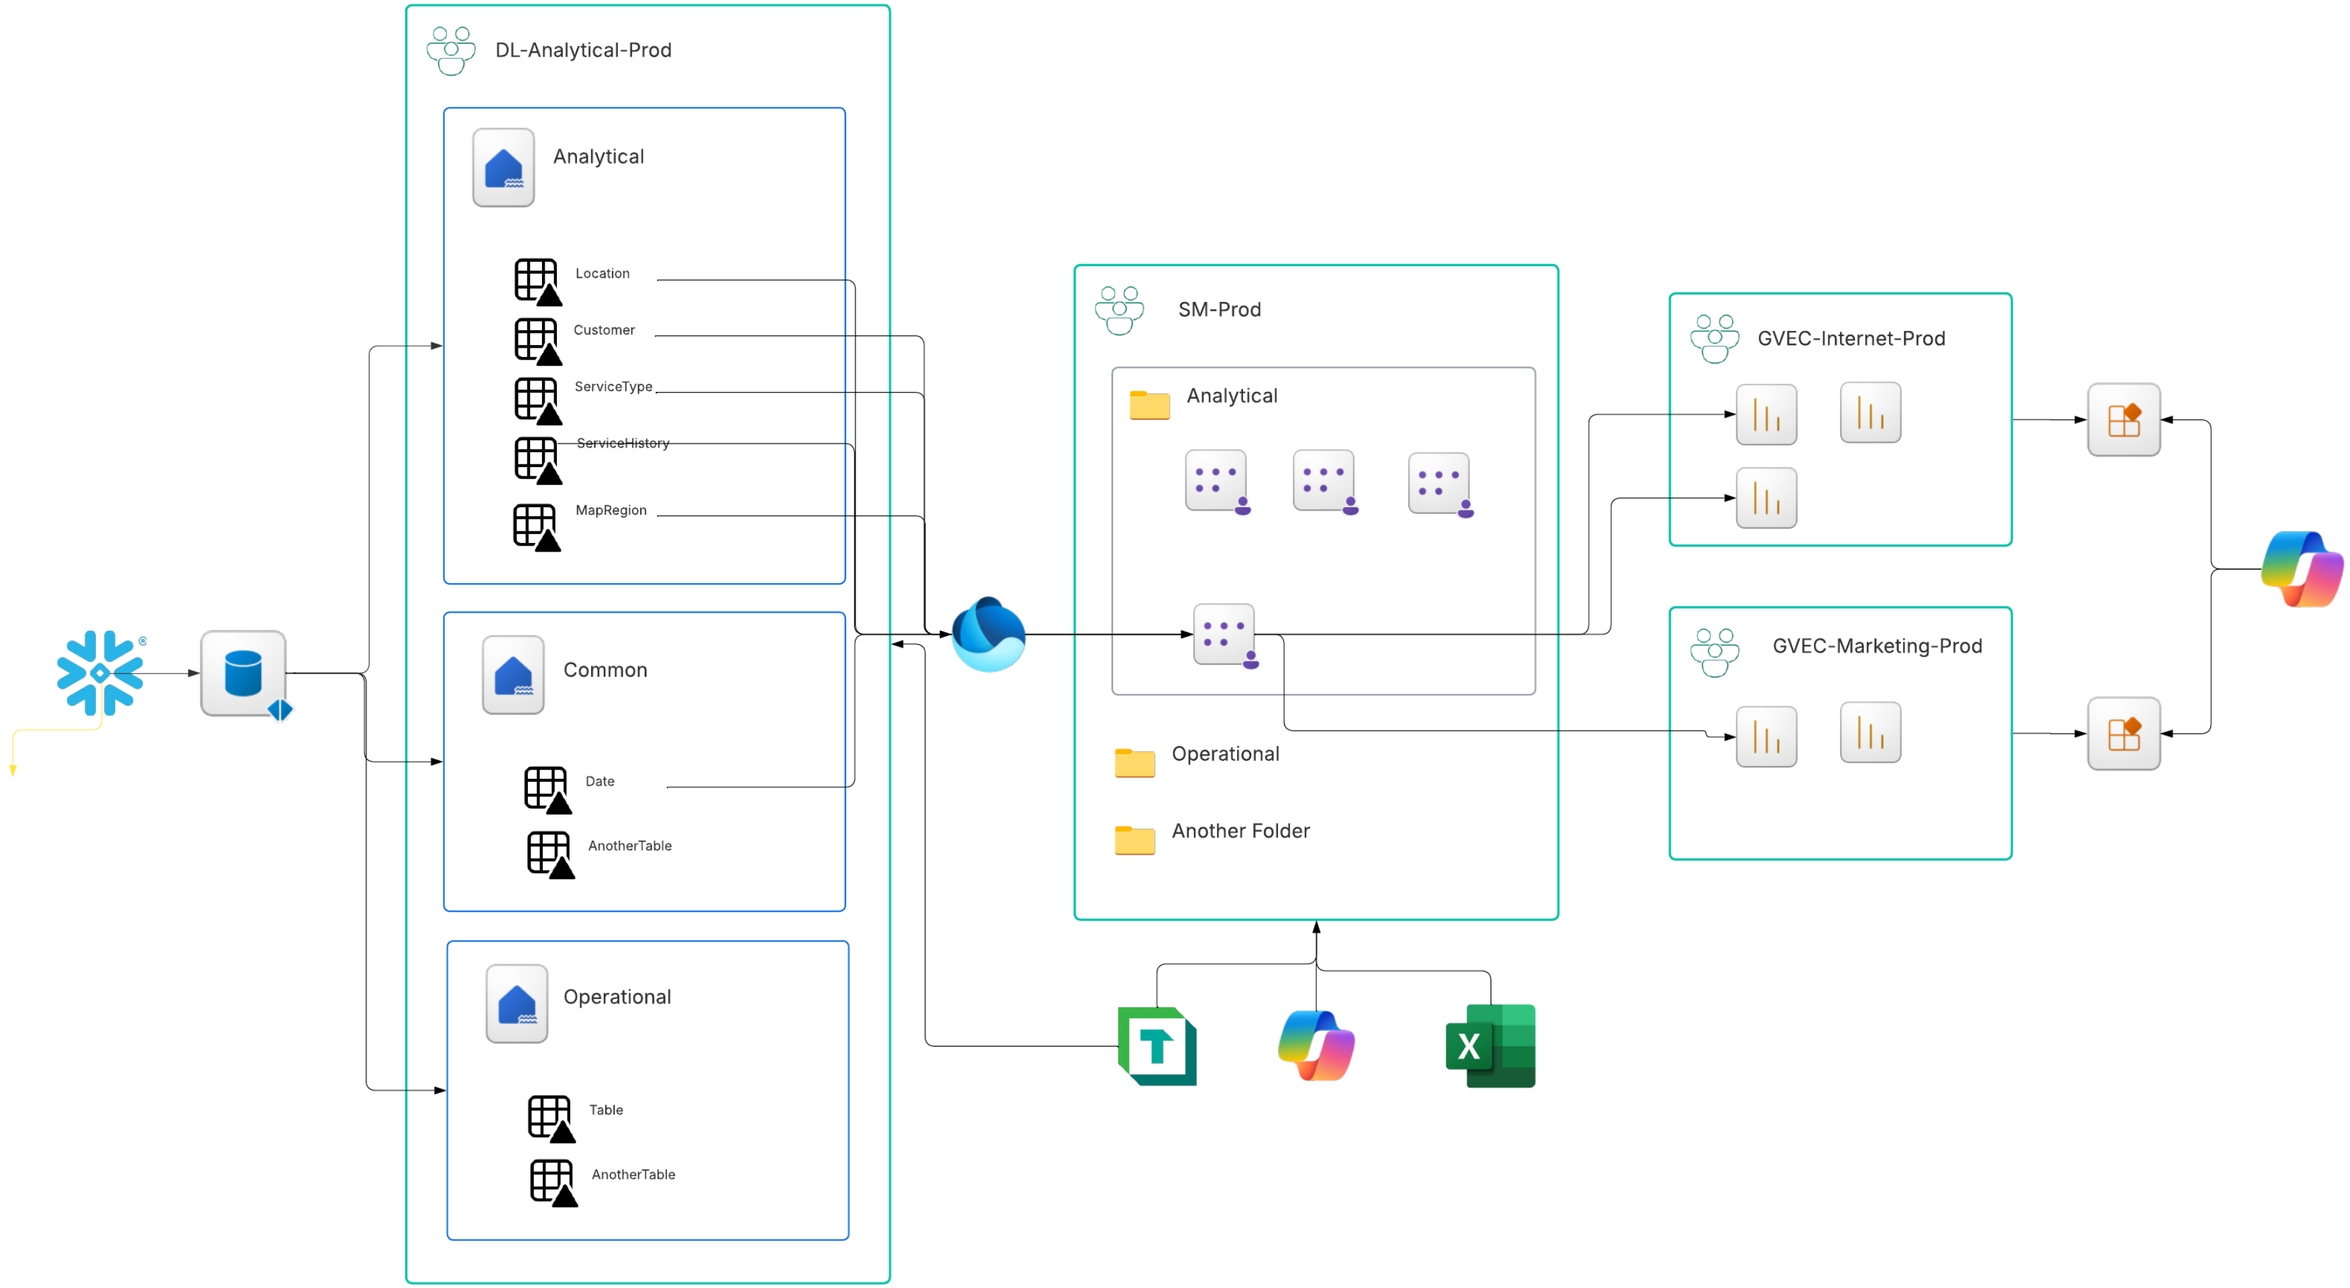
DL-Analytical-Prod
Analytical
Location
Customer
ServiceType
ServiceHistory
MapRegion
Common
Date
AnotherTable
Operational
Table
AnotherTable
SM-Prod
Analytical
Operational
Another Folder
GVEC-Internet-Prod
GVEC-Marketing-Prod

A dedicated workspace will be used to land data into a Lakehouse with a well-organized structure. Semantic models will connect to this single source of truth via DirectLake, eliminating the need for scheduled refreshes. There will only be one production workspace, and all business logic transformations will be mirrored from our Snowflake warehouse

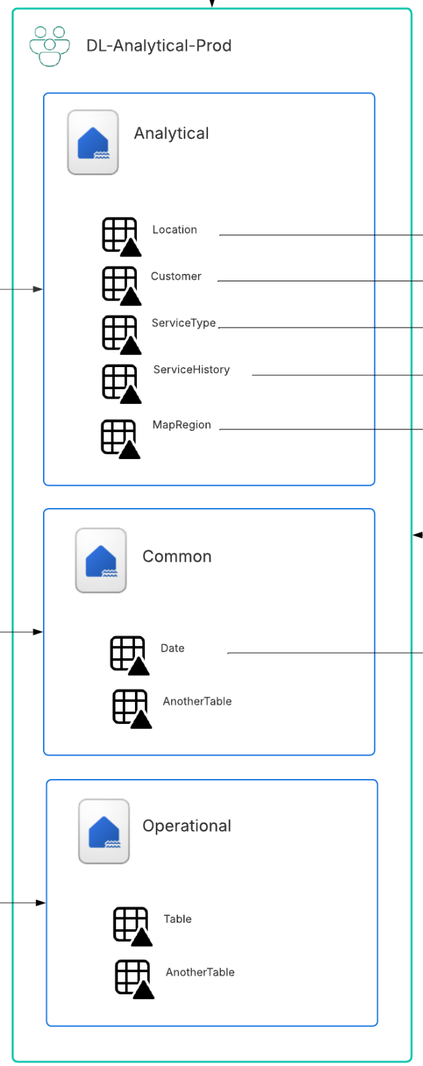

Shared semantic models will connect to the data lake using DirectLake. This approach combines the advantages of both Import and DirectQuery modes. A key benefit is that scheduled refreshes for semantic models and reports will no longer be required.

A dedicated workspace will be created specifically for semantic models, featuring a defined folder structure for easy organization and discoverability. This layer will also manage access permissions for each model. If sensitive semantic models need to be siloed, a separate workspace will be established to handle that data securely

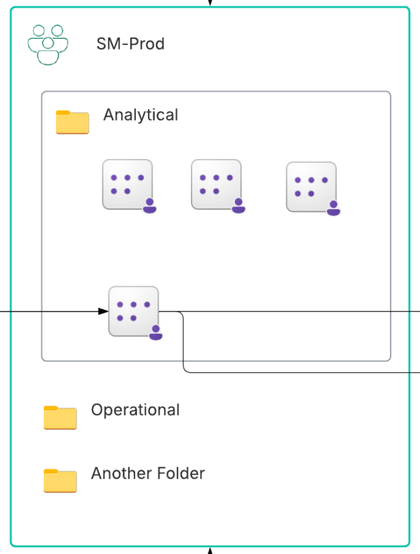

The existing production workspaces and apps will remain unchanged, with only 'Prod' added to the end of each workspace name. As reports are repointed to shared semantic models, these new reports will gradually replace the current ones. End users should experience no disruption during this transition.

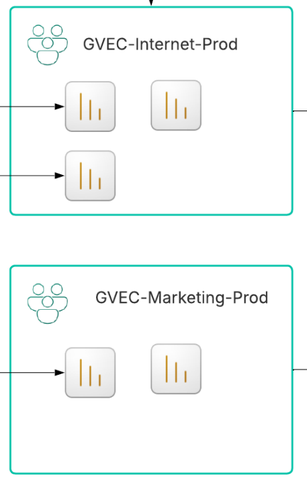

Each report will include an AI-enhanced semantic model, enabling users to leverage Copilot within the app to ask questions in natural language.

Case Study A

Analytical tables from Snowflake are mirrored or imported into Fabric, possibly eliminating the need for scheduled refreshes thanks to continuous syncing.

Tabular Editor will be the primary tool for creating and maintaining semantic models. It is a Microsoft-approved external tool and is included in the Fabric Engineering and Fabric Solution Architecture certifications.

Semantic models will be AI-prepped for use across Copilot experiences, including Power BI, Microsoft 365 Copilot, agents, and other integrated tools.

To democratize our data, end users will be able to connect to certified semantic models that include validated business logic and governance, enabling them to explore data directly in Excel.