

TAKEAWAYS FROM FABCON EUROPE 2026

Easier To Build, Harder To Govern

What FabCon Europe 2026 Tells Us About Where BI Is Heading

Barcelona, 28 September to 1 October 2026. What was announced, the industry shift it reflects, and the settings to check this week.

• **Current as of 5 October 2026**

Short shelf life. Settings and preview status may change after this date.

01 What shipped

GENERALLY AVAILABLE

- GA** **Fabric IQ in Copilot Chat**
Shared intelligence layer over OneLake, semantic models and ontologies.
- GA** **Fabric IQ MCP server**
External MCP clients can reach Fabric IQ.
- GA** **Restrict from Copilot**
Admins choose which models participate in AI experiences.
- GA** **Per-model read-only user access**
Editor-managed setting on each semantic model. Enabled by default.
- GA** **Sensitivity labels and DLP into Copilot**
Labels and data loss prevention policies apply inside Copilot.
- GA** **Capacity overage, surge protection, F4096, F8192**
Overage handling, surge controls and two larger capacity SKUs.

GA

Generally available

PREVIEW

Public preview

PREVIEW

- PREVIEW** **Fabric Apps in Power BI Pro**
Natural-language generation of data apps with writeback.
- PREVIEW** **Microsoft 365 Copilot in Power BI**
Model-backed answers inside Microsoft 365 apps.
- PREVIEW** **IQ Sharing**
Sharing governed data and context with customers and partners.

ANNOUNCED

- COMING SOON** **F0 and on-demand billing**
Pricing not yet published.

TWO SIDES OF THE SAME SHIFT

The distance from question to answer got shorter. The number of ways to reach a model went up.

The announcements at FabCon Europe reflect a broader move across BI: AI is becoming a standard way to reach governed data. Both columns below read the same set of features from two seats: the person building with them, and the person accountable for who can reach the data behind them.

02 What got easier

One intelligence layer over your data

Fabric IQ in Copilot Chat puts OneLake, semantic models and ontologies behind one conversational surface. Answers draw on the models your team already built and maintains.

Insight in the flow of work

Microsoft 365 Copilot in Power BI brings model-backed answers into the apps people already have open. Nobody has to find the right report first.

Data apps without a development cycle

Fabric Apps in Power BI Pro generates data apps with writeback from a natural-language description, on the Pro licence teams already hold.

An open door for agents

The Fabric IQ MCP server lets any MCP-compatible client query Fabric IQ through a standard protocol, without a custom integration.

Governance that travels with the data

Sensitivity labels and DLP now apply inside Copilot, and Restrict from Copilot gives admins a per-model switch for AI participation.

03 What got harder

Four access paths into one model

A semantic model once reached through reports can now be reached through Power BI Copilot, Microsoft 365 Copilot, data agents and MCP servers. Each path is a different client in front of the same data.

Cost predictability

AI workloads from every path draw on the same capacity. Overage and surge protection change how spikes are absorbed and billed. F0 and on-demand billing are announced, with pricing not yet published.

Attribution across product boundaries


When the question is asked in one product and answered from a model in another, showing which measure and which filter produced a number takes more work to trace.

The external boundary

Microsoft's FabCon materials do not address guest or external access to models through Copilot, data agents or MCP. IQ Sharing is the announced path for customers and partners, and it is in preview.

04 Enabled by default

Read-only user access to semantic models through Power BI Copilot, Microsoft 365 Copilot, data agents and MCP servers is governed by a per-model setting. Model editors manage it, and it is **enabled by default**.



PATH	WHAT IT EXPOSES	WHO REACHES IT	WHAT TO CHECK
1 · Power BI Copilot	Model data through natural-language questions in Power BI	Users with read access to the model	Whether row-level security covers questions people can now ask, not only the visuals in reports
2 · Microsoft 365 Copilot	Model-backed answers in Microsoft 365 apps	Model readers, working outside Power BI	Sensitivity labels on the models you would least want summarised in a chat or a document
3 · Data agents	Model data through agents configured on top of it	Users of each agent who hold read access	Which agents point at which models, and who owns each agent
4 · MCP servers	Model access for MCP clients, including tools outside Microsoft's stack	Readers connecting through any MCP-compatible client	Which MCP clients are in use against your tenant, and which models they can reach

05 The boundary

AS OF 5 OCT 2026

IQ Sharing, the announced path for sharing governed data and context with customers and partners, is **PREVIEW**. Its permission granularity is not yet documented. Until it is, treat external delivery as a separate decision from internal AI enablement.

06 What to do this week

From *Easier To Build, Harder To Govern: Key Takeaways From FabCon Europe 2026*

Current as of 5 Oct 2026

- ☐ **List every semantic model with read-only user access on.** It is enabled by default, so start from the assumption that all of them are.
- ☐ **Apply Restrict from Copilot** to models that should not take part in AI experiences yet.
- ☐ **Confirm sensitivity labels** on models that carry regulated or client data, so DLP policies follow that data into Copilot.
- ☐ **Inventory data agents and MCP clients** already pointed at your models, and name an owner for each.
- ☐ **Review capacity overage and surge protection settings** before AI usage grows across all four paths.
- ☐ **Decide your position on external users** before anyone asks. IQ Sharing is in preview and its permission model is not yet documented.

Next: 7 Questions To Answer Before External Users Touch AI-Generated Insight

thereportinghub.com

07 · WHERE REPORTING HUB FITS

When the audience is outside your organisation, **the AI has to show its work.**

Reporting Hub is governed delivery for Power BI: a branded portal that runs inside your Azure tenant, with no data egress. Its BI Genius AI works to an admin-configured scope, one agent per audience, with source attribution, DAX transparency and an audit trail behind every answer. It is built on Microsoft Foundry and sits alongside the Fabric and Copilot capabilities described in this paper.

**Microsoft Solution
Accelerator Partner**

4.5 / 5 on G2
Verified reviews, SMB to Enterprise

300,000+ active users

Deployed in 30+ countries

Start a 30-day free trial

No credit card required. BI Genius Starter is included on every plan.

Start Now →

SOURCES

Microsoft: FabCon and SQLCon 2026 in Barcelona

FabCon Europe 2026: The Fabric announcements that matter

Microsoft Fabric blog: Bringing governed analytics into the flow of work

Schellman 2026 AI governance research