Thank you sponsors!



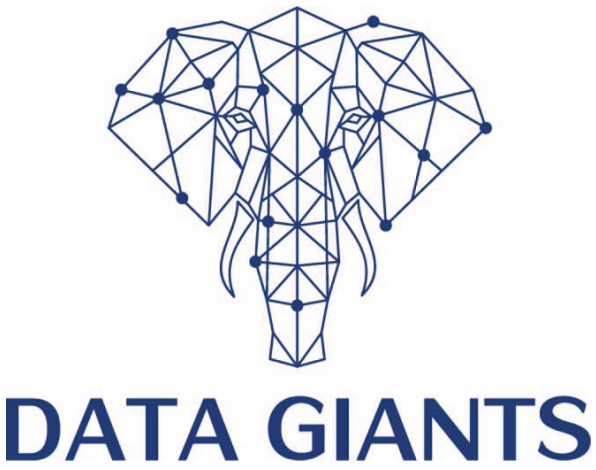
























Are your semantic models ready for Copilot?

Lars Andersen
Microsoft
Fabric CAT



Thank you sponsors!

INSPARK

twintos

ilionx

InfoSupport
Solid Innovator

redgate

sopra steria

DATA GIANTS

BRANDER
company

SQL brander.nl
DATA brander.nl
CRM brander.nl
ERP brander.nl

Bull
Master your AI destiny

dbWatch

BBTG

Data Saturdays

Agenda

- Data answering in Fabric
- Enabling Copilot for Power BI
- Understand how to use Copilot
- Grounding data for Copilot
- Best practice and prepare semantic models for Copilot

Data answering in Fabric



Data Answering

Business user AI experiences in Fabric



Copilot

Generalist



Fabric Data Agents

Specialist

Similarities

You may have noticed some key similarities



Copilot



Data Agents

- ✓ Highly conversational
- ✓ Answer user questions about their data
- ✓ Human in the loop setup for best results

Different Goals

Built to help customers get data insights using AI



Copilot

Extract **more** value & new insights from **existing data** content & context



Data Agents

Craft **new** agents to deliver a **tailored experience** on top of OneLake

Scope

What data can be reasoned over?



Copilot

- Visual layer: All Reports
- Data layer: Semantic models

- *All* content the user has access to PBI



Data Agents

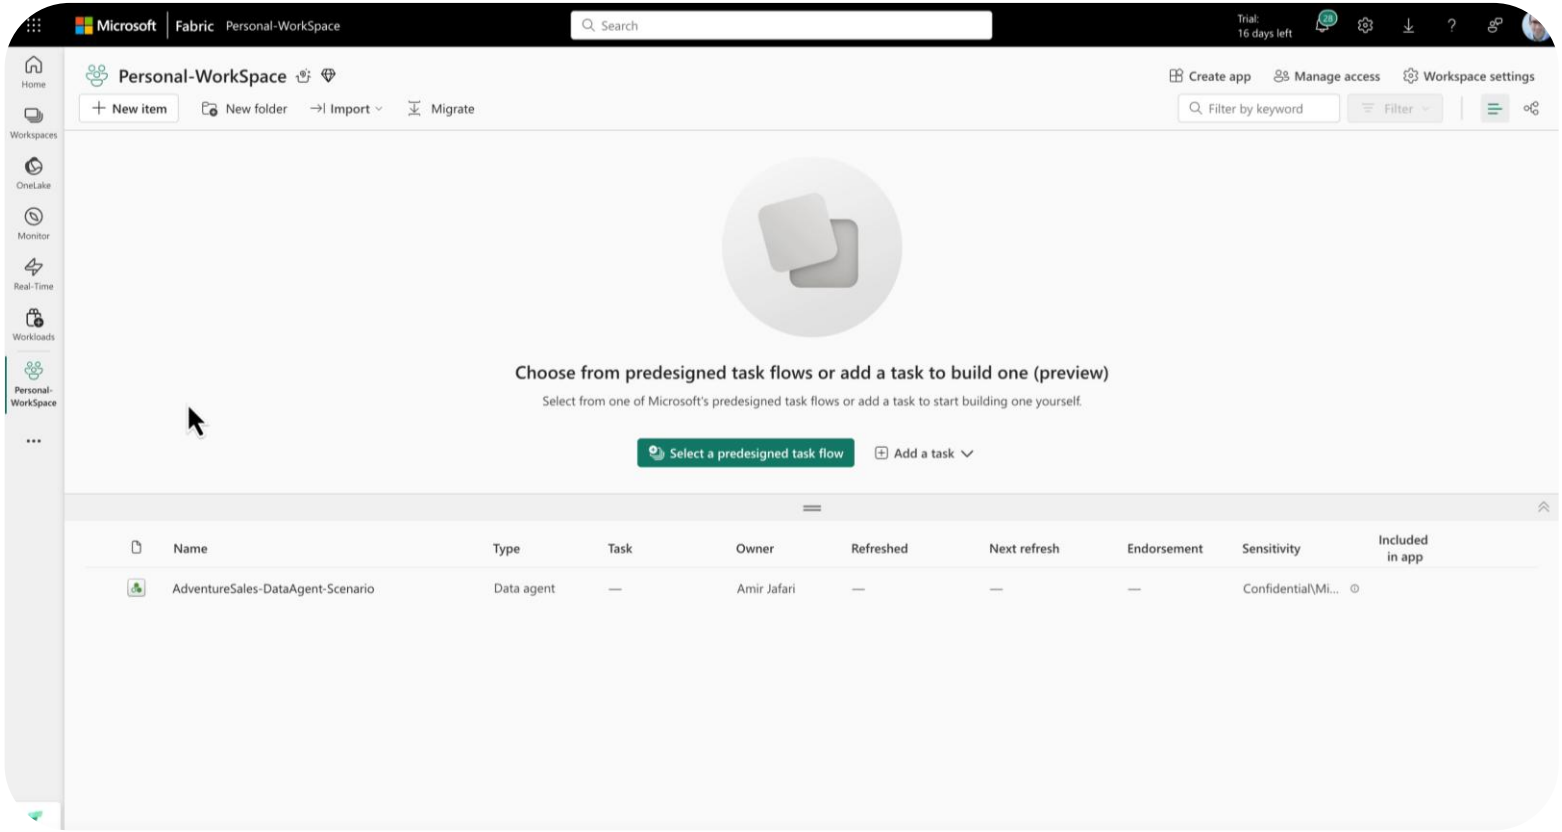
- **Only** data layer: Structured & unstructured, OneLake & mirrored sources, ontologies, semantic models

- Just what is added to the data agent
- Limit to 5 data sources (SDK not limited)

Fabric Data Agent

Enrich custom copilots with Fabric data expertise

Generally Available



Create a powerful data expert **tailored to reason** over and synthesize data in OneLake for insights.

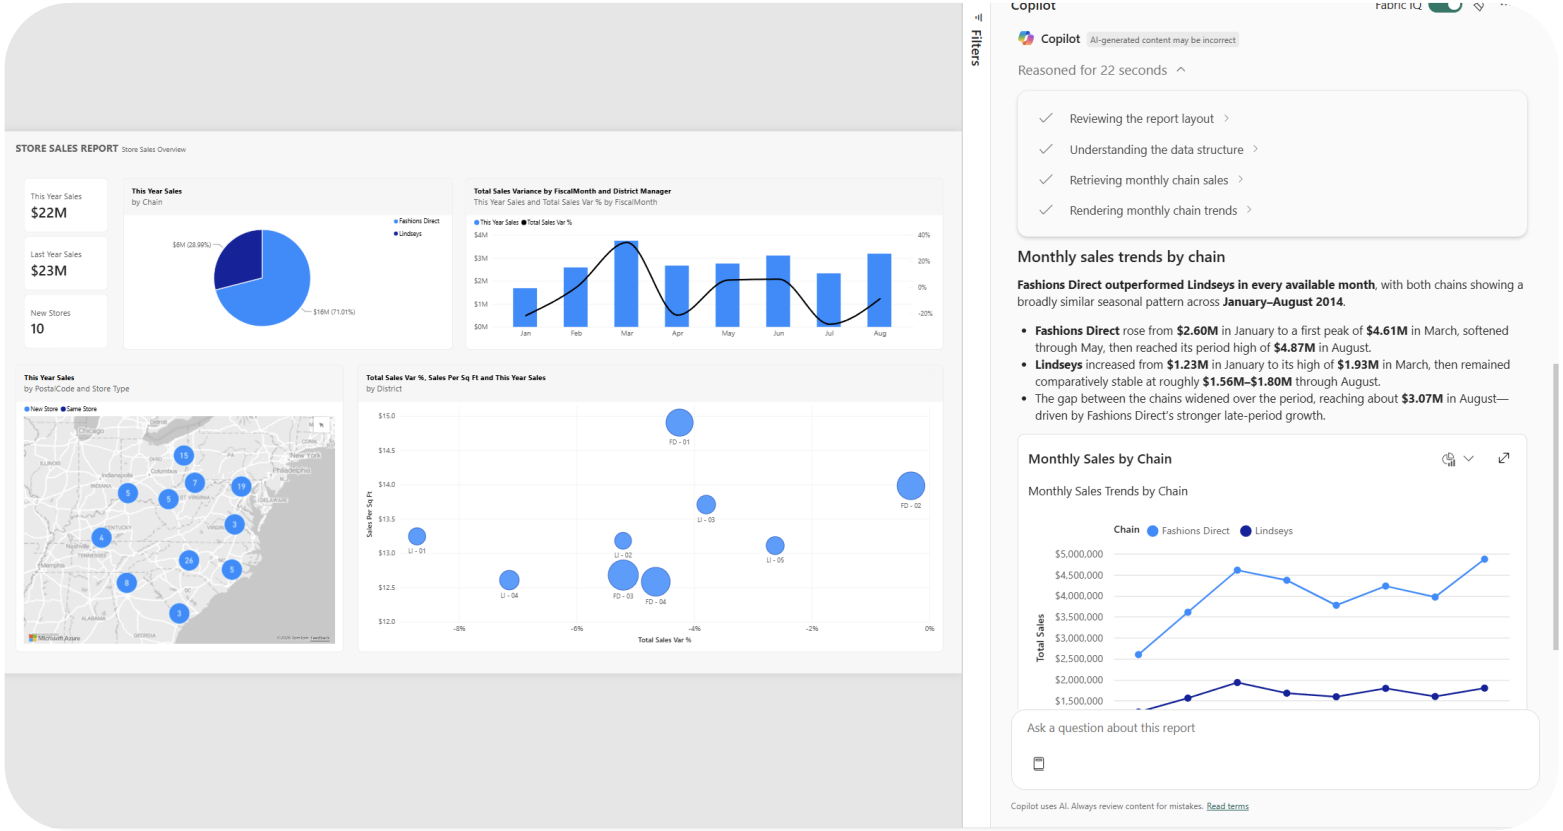
Leverage **structured & unstructured data sources** to provide holistic answers.

Security permissions, including row and column level security, **are respected**.

Generally Available

Copilot in Power BI report

Get answers from report and semantic model



Fabric IQ adds reasoning transparency.

Copilot can generate **new analyses,** not just read existing visuals

Data quality and semantic model **preparation** matter

Fabric IQ with Power BI

Business context for every agent, now generally available

GA

Fabric IQ in Copilot Chat

Get quick answers

GA

Fabric IQ plugin in Copilot Cowork

Ask in the flow of work

GA

Fabric IQ MCP server

Build with agents

Bring trusted Power BI insights to your users

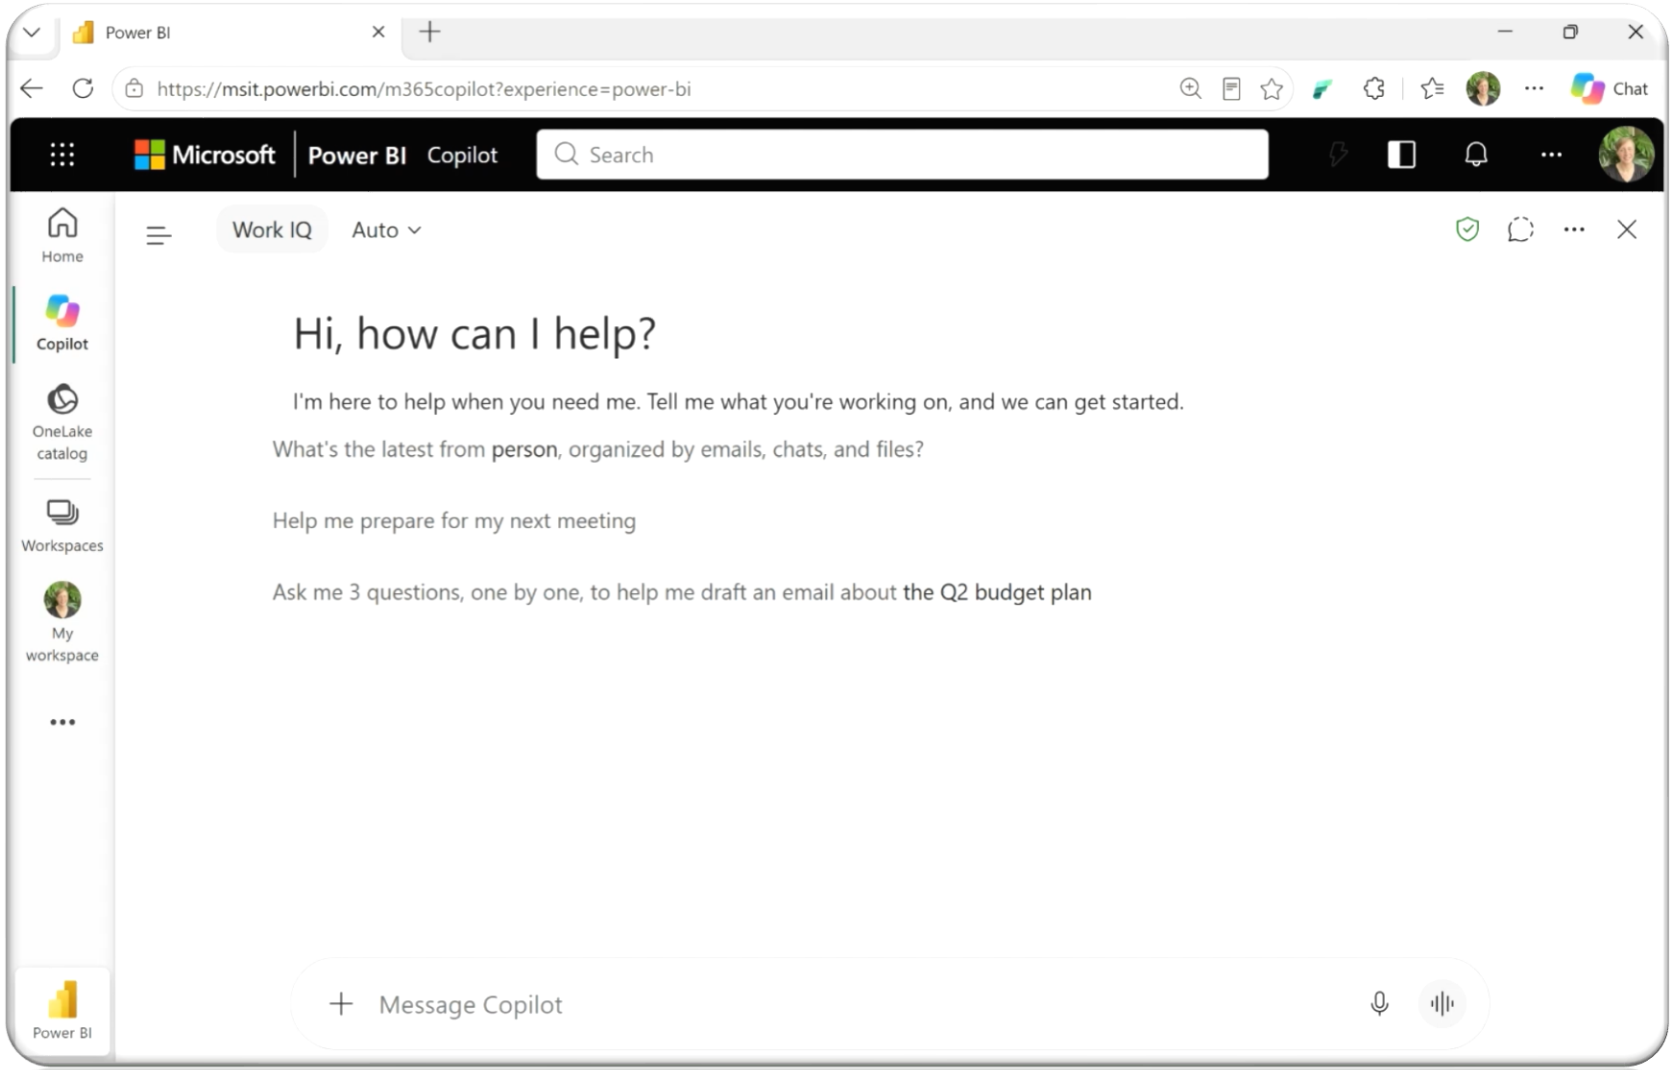
Copilot in Fabric and Power BI

Public Preview

Option to use the full **Microsoft Copilot** directly from Power BI – instead of PBI standalone Copilot

Find recent **Power BI data and Work IQ**, without ever leaving Fabric

Kick off background tasks with **Cowork** in the flow of work



Comparison

FEATURE	Copilot	Copilot Cowork	Data agent
Ground a Power BI report	Fully supported	Fully supported	Partly supported
Attach / paste / name a report	Fully supported	Fully supported	Not supported
Lakehouse / Warehouse / KQL sources	Not supported	Not supported	Fully supported
Custom instructions and examples	Partly supported	Partly supported	Fully supported
Chain into email, docs, or meetings	Partly supported	Fully supported	Not supported
User permissions and RLS	Fully supported	Fully supported	Fully supported
Reusable domain-specific expert	Not supported	Partly supported	Fully supported

Demo: Copilot, Cowork, and Data agent

Scope for the rest of this session: Grounding semantic model for AI

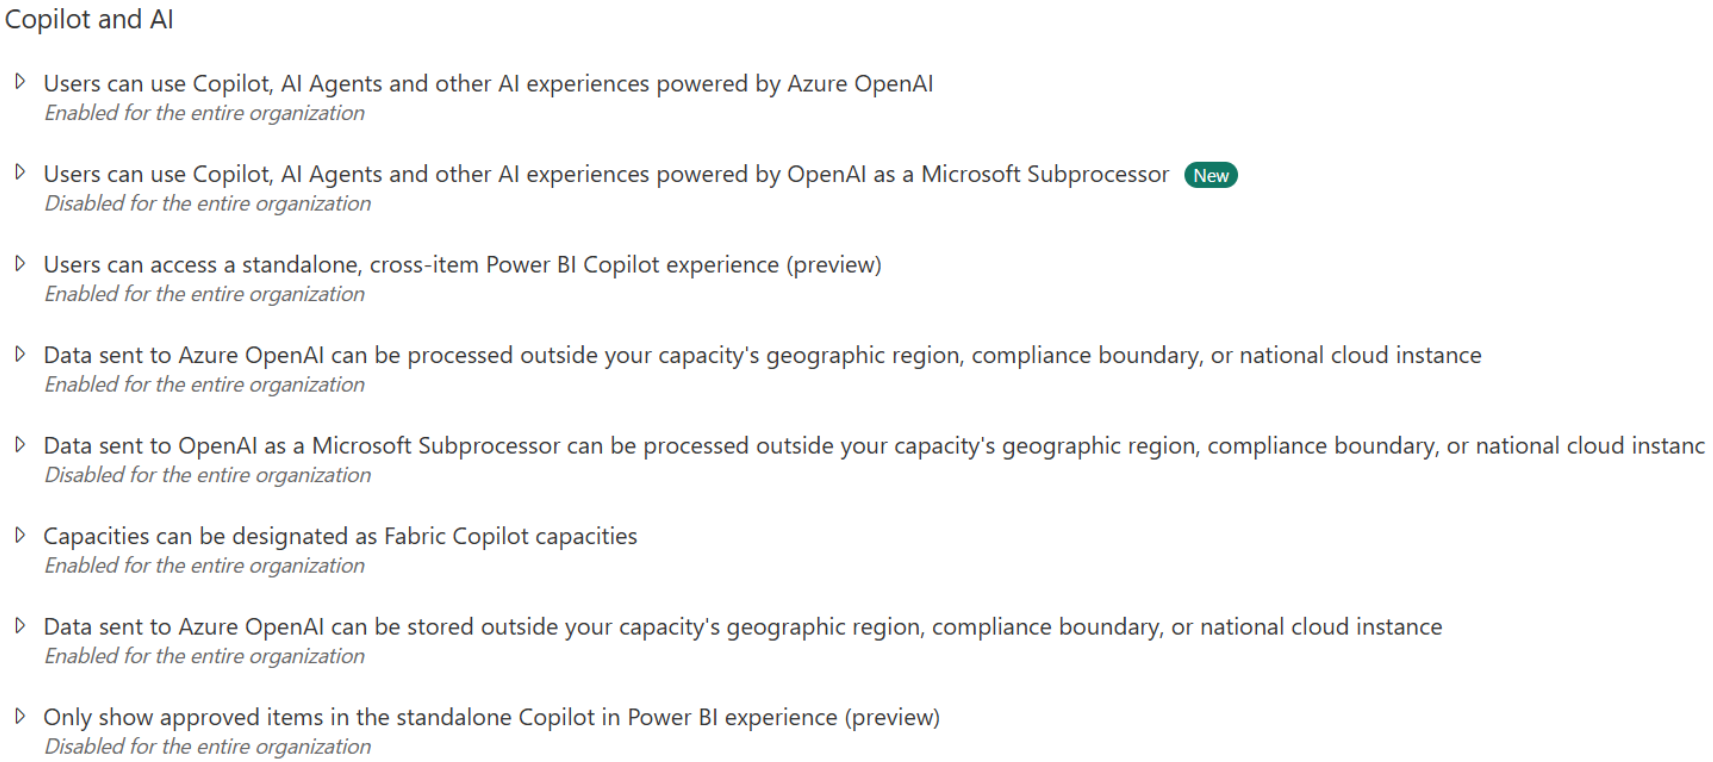


Enabling Copilot for Power BI



Tenant settings

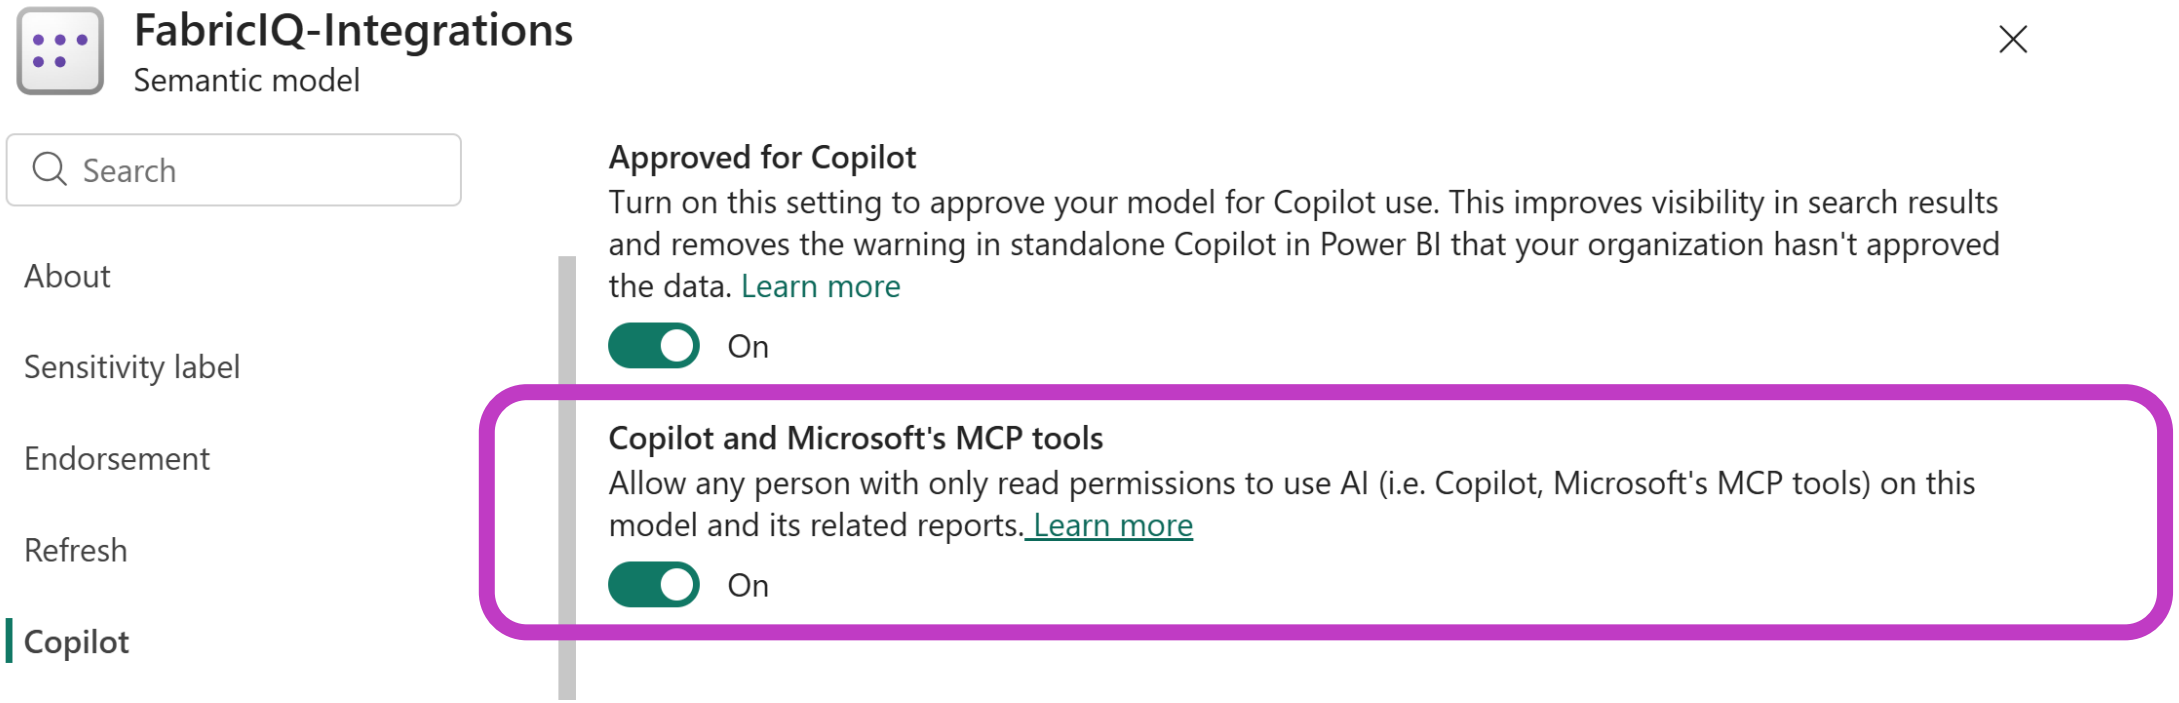
- Enabling of Copilot in Power BI is documented here Enable Fabric Copilot for Power BI
- As of today, these tenant settings exist:



Tenant settings

- *Users can use Copilot and other features powered by Azure OpenAI*: **Must be enabled to use Copilot**
- *Users can use Copilot, AI Agents and other AI experiences powered by OpenAI as a Microsoft Subprocessor*: **Optional**
- *Users can access a standalone, cross-item Power BI Copilot experience (preview):* **Optional**
- *Data sent to Azure OpenAI can be processed outside your capacity's geographic region, compliance boundary, or national cloud instance*: **Required if Azure OpenAI isn't available in capacity region**
- *Data sent to OpenAI as a Microsoft Subprocessor can be processed outside your capacity's geographic region, compliance boundary, or national cloud instance*: **Optional**
- *Capacities can be designated as Fabric Copilot capacities*: **Optional**
- *Data sent to Azure OpenAI can be stored outside your capacity's geographic region, compliance boundary, or national cloud instance*: **Optional, applicable for Copilot in Notebooks and Data agent feature**
- *Only show approved items in the standalone Copilot in Power BI, and in the Power BI Agent experiences (preview)*: **Optional**

Creator control over semantic model use in AI



Approved for Copilot still improves visibility in search results

Now, creators can **disable** Copilot & MCP tool access
to a semantic model and its reports – also changeable via batch API

Understand how to use Copilot



Understand how to use Copilot

Understand limitations

Report > Model

Supported v. unsupported query types

Perfecting questions

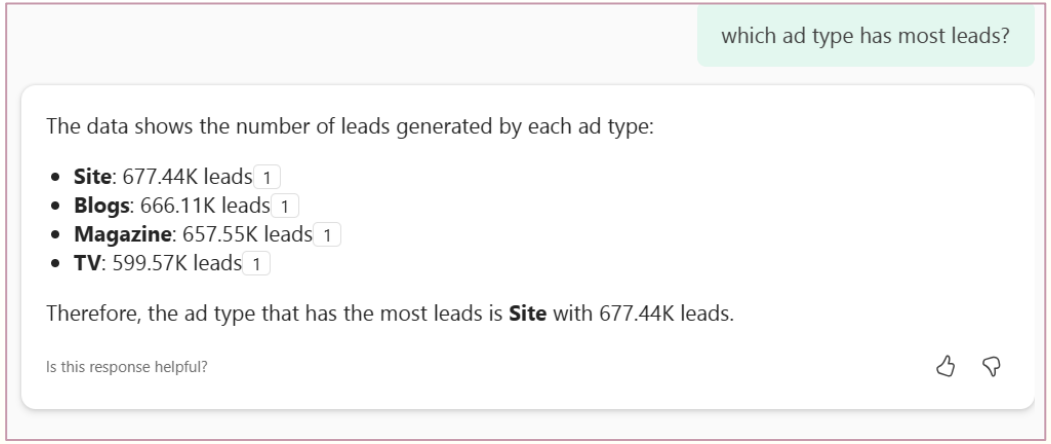
Ideal sentence structure

Understand how to use Copilot in report

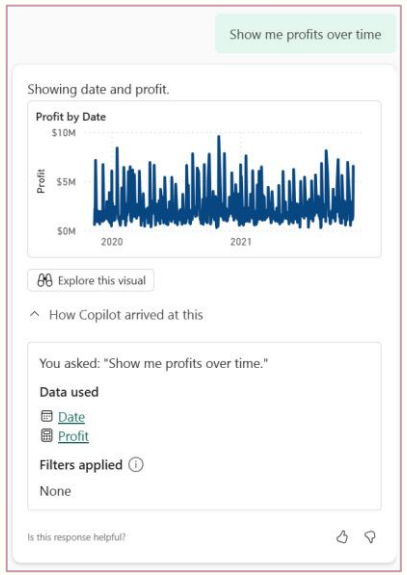
Copilot limitations

Report visuals are prioritized over the model information

Report response



Model response



Understand how to use Copilot

Copilot limitations

- Columns and fields in the report or data model

“What were the top 5 selling products in North America last month?”

where product, region, and date are already columns in the data, and Copilot helps with top N filtering to produce the visual.

- Explanations, summaries, or “show me” questions

“Tell me the average price of gasoline per gallon over the last 30 days.”

where price per gallon is a field already contained in the model, and Copilot helps by taking the average and filtering by relative date to produce the visual.

Understand how to use Copilot

Copilot limitations

- ✗ Why questions
- ✗ Diagnostic analysis like root cause or forecasting

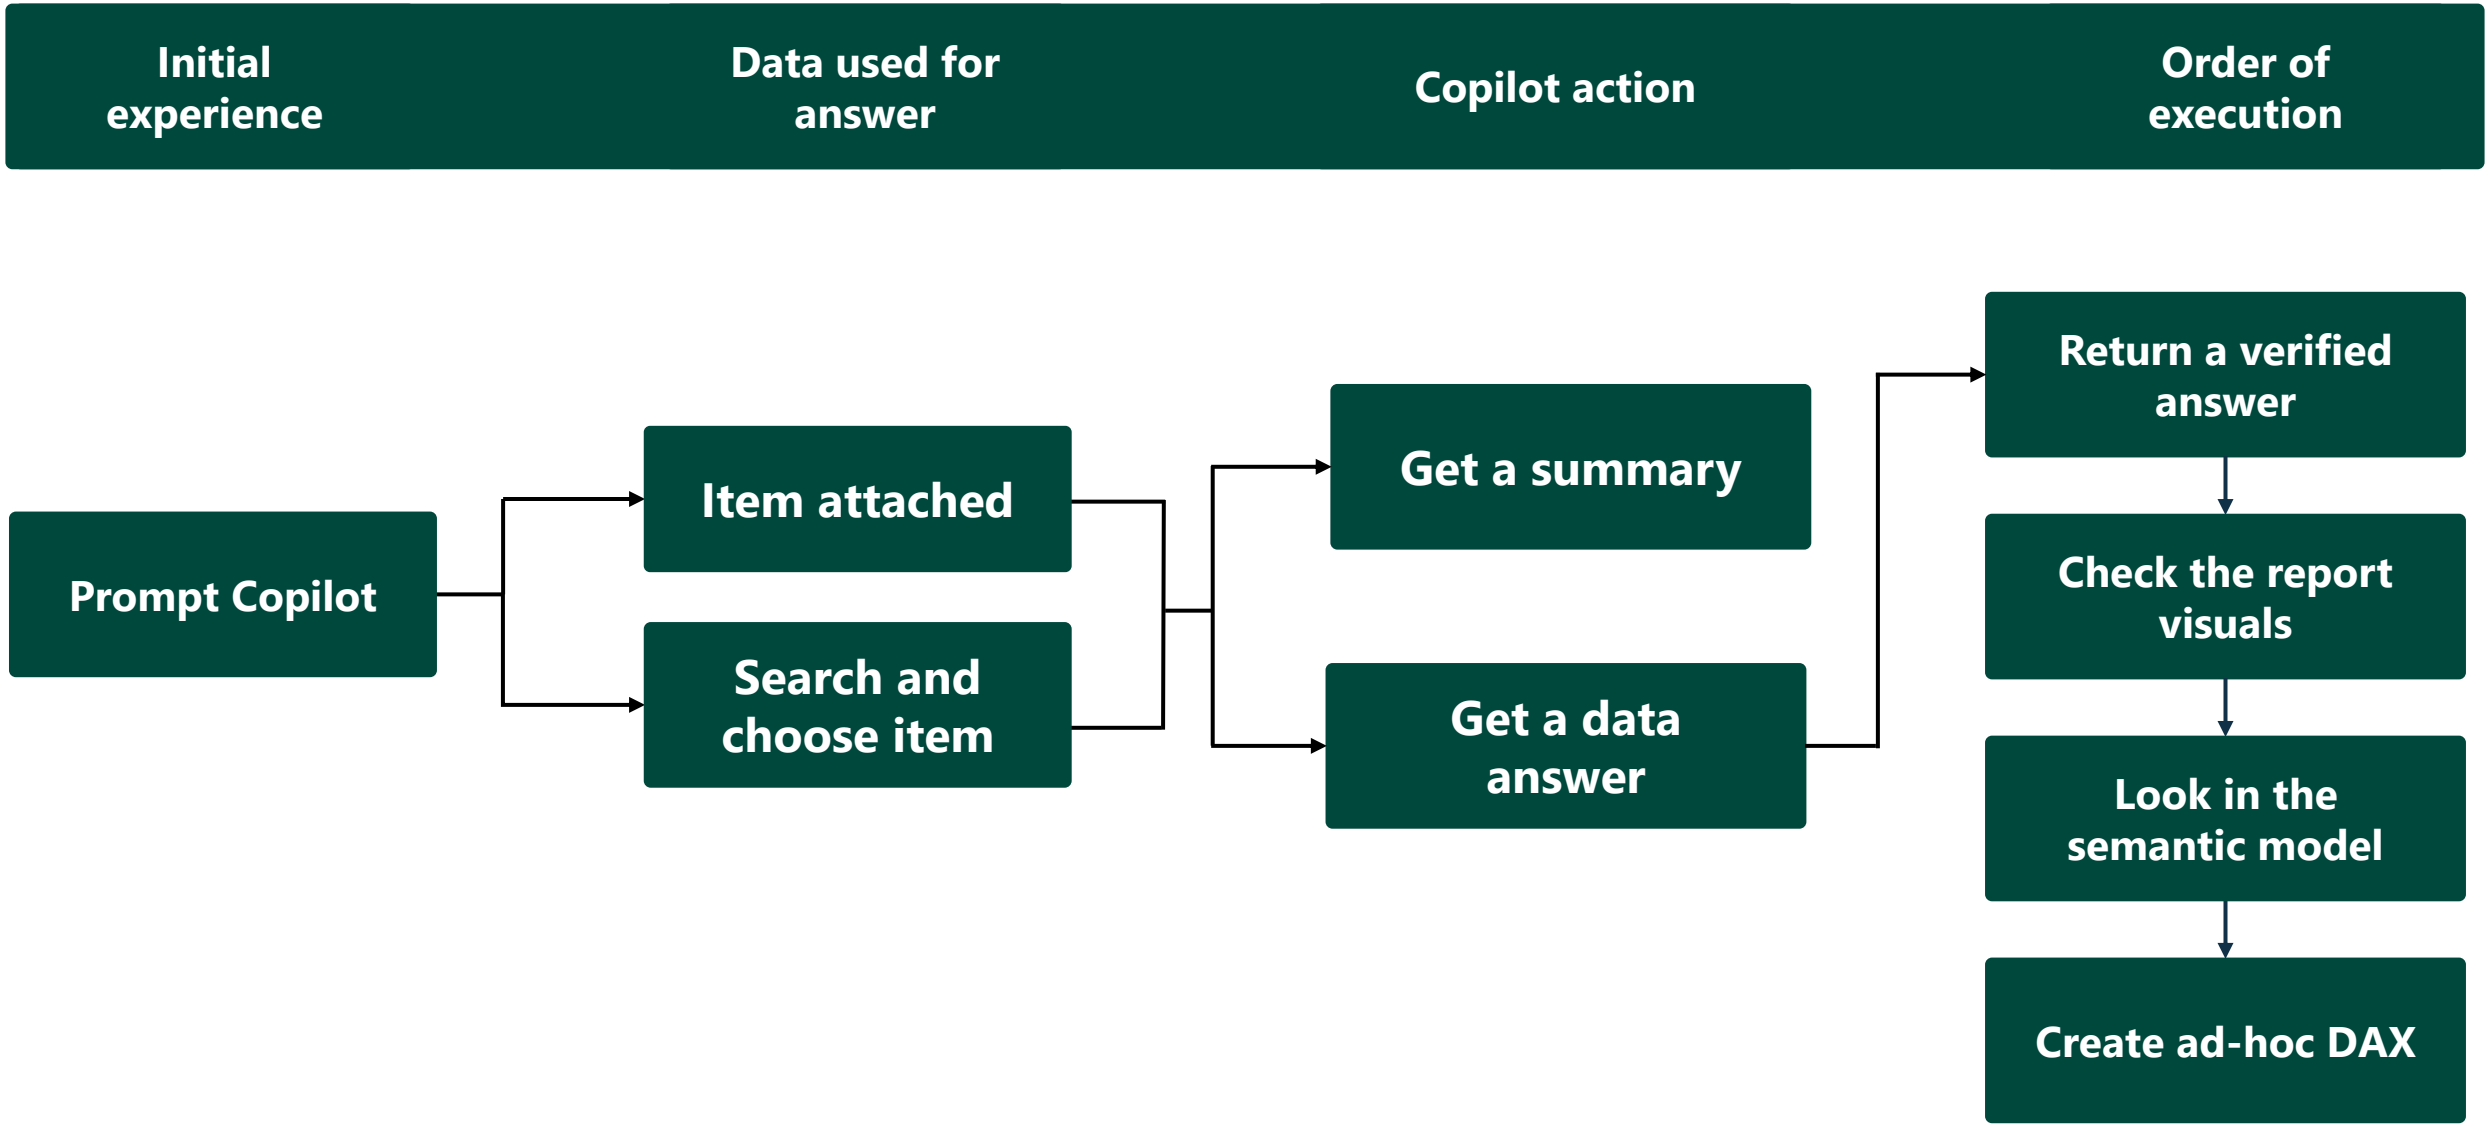
"How many books do you think we will sell next year?"

This question asks for forecasting, which isn't currently supported

"Why do our sales go down every July?

This question involves generating deeper insights beyond the provided data.

Behind the scenes (simplified) **in report**



Understand how to use Copilot

Perfecting questions

Ideal sentence structure

Average ticket sales by movie genre
... grouped by region
... released in 2024

Top 5 movie titles by total ticket sales
... as a column chart

1. preferred sort order
2. aggregation on value/measure by dimension field
3. legend
4. filters
5. visual type

Grounding data for Copilot



Configuring and grounding data for Power BI Copilot: **Prep data for AI (preview) – available now…**

Verified Answers

Ensure **quality** and **consistency** of responses for the most critical business topics.

Schema Selection

Limit the scope of data that Copilot can answer questions about

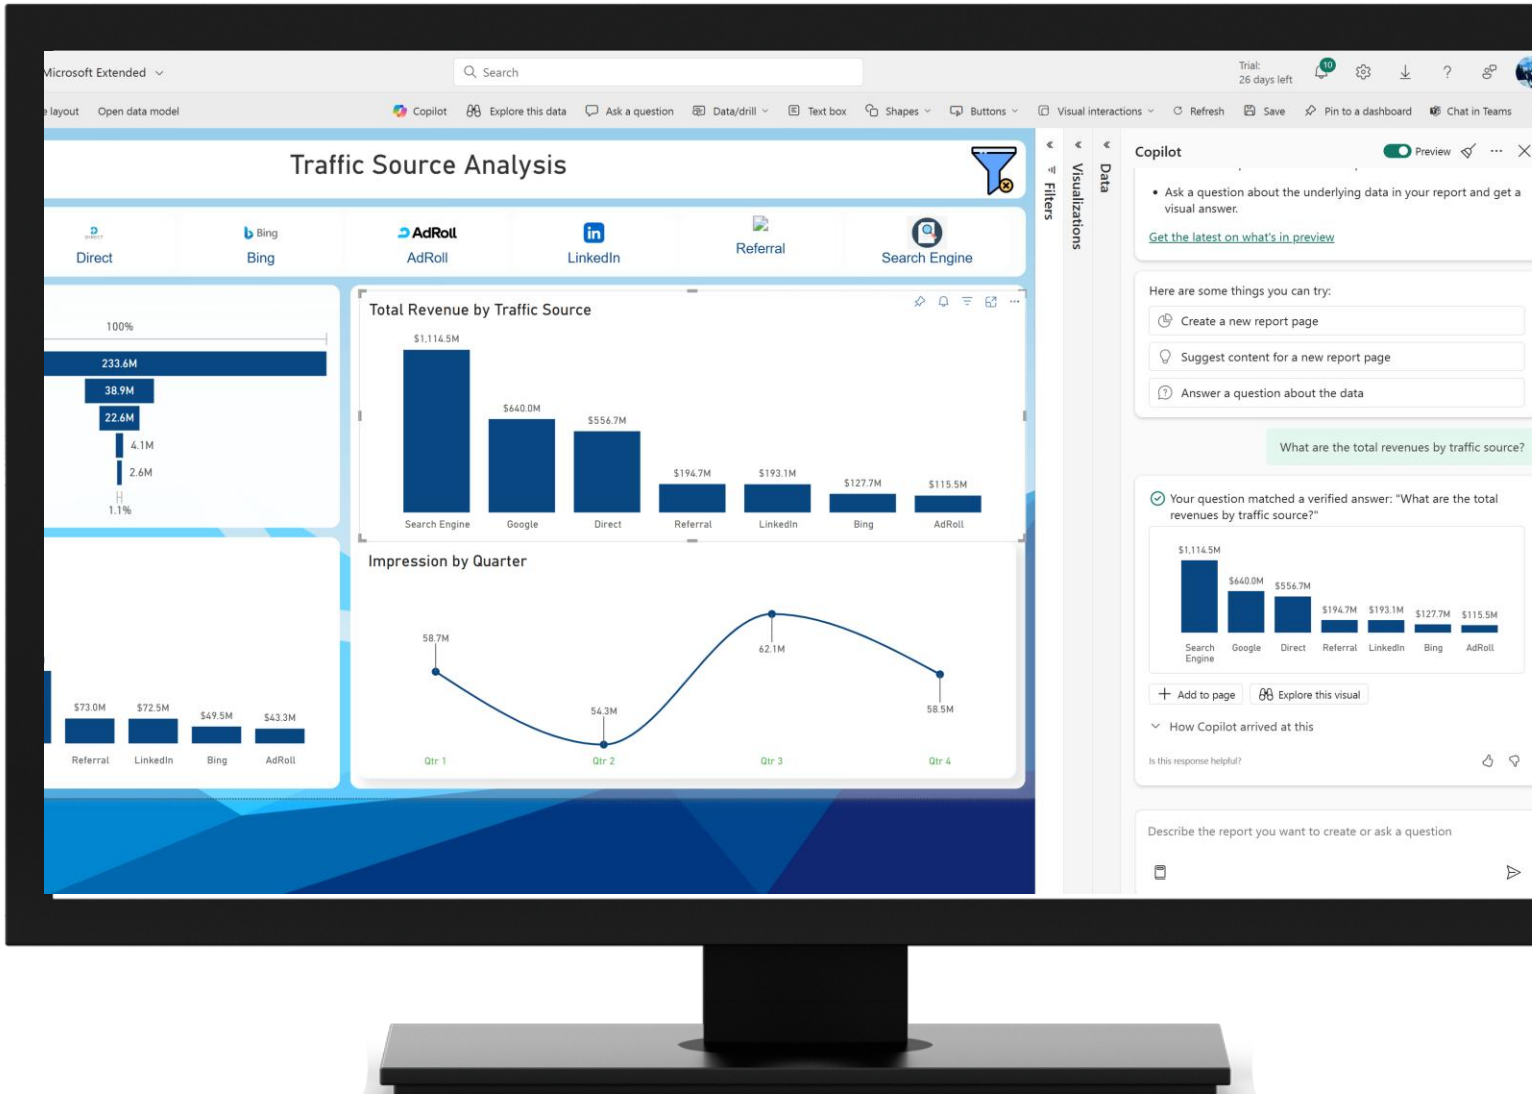
Custom Instructions

Provide business context and help Copilot guide how to talk about your data

Power BI Copilot grounding: Verified Answers

Verified Answers

Ensure quality and consistency of responses for the most critical business topics.

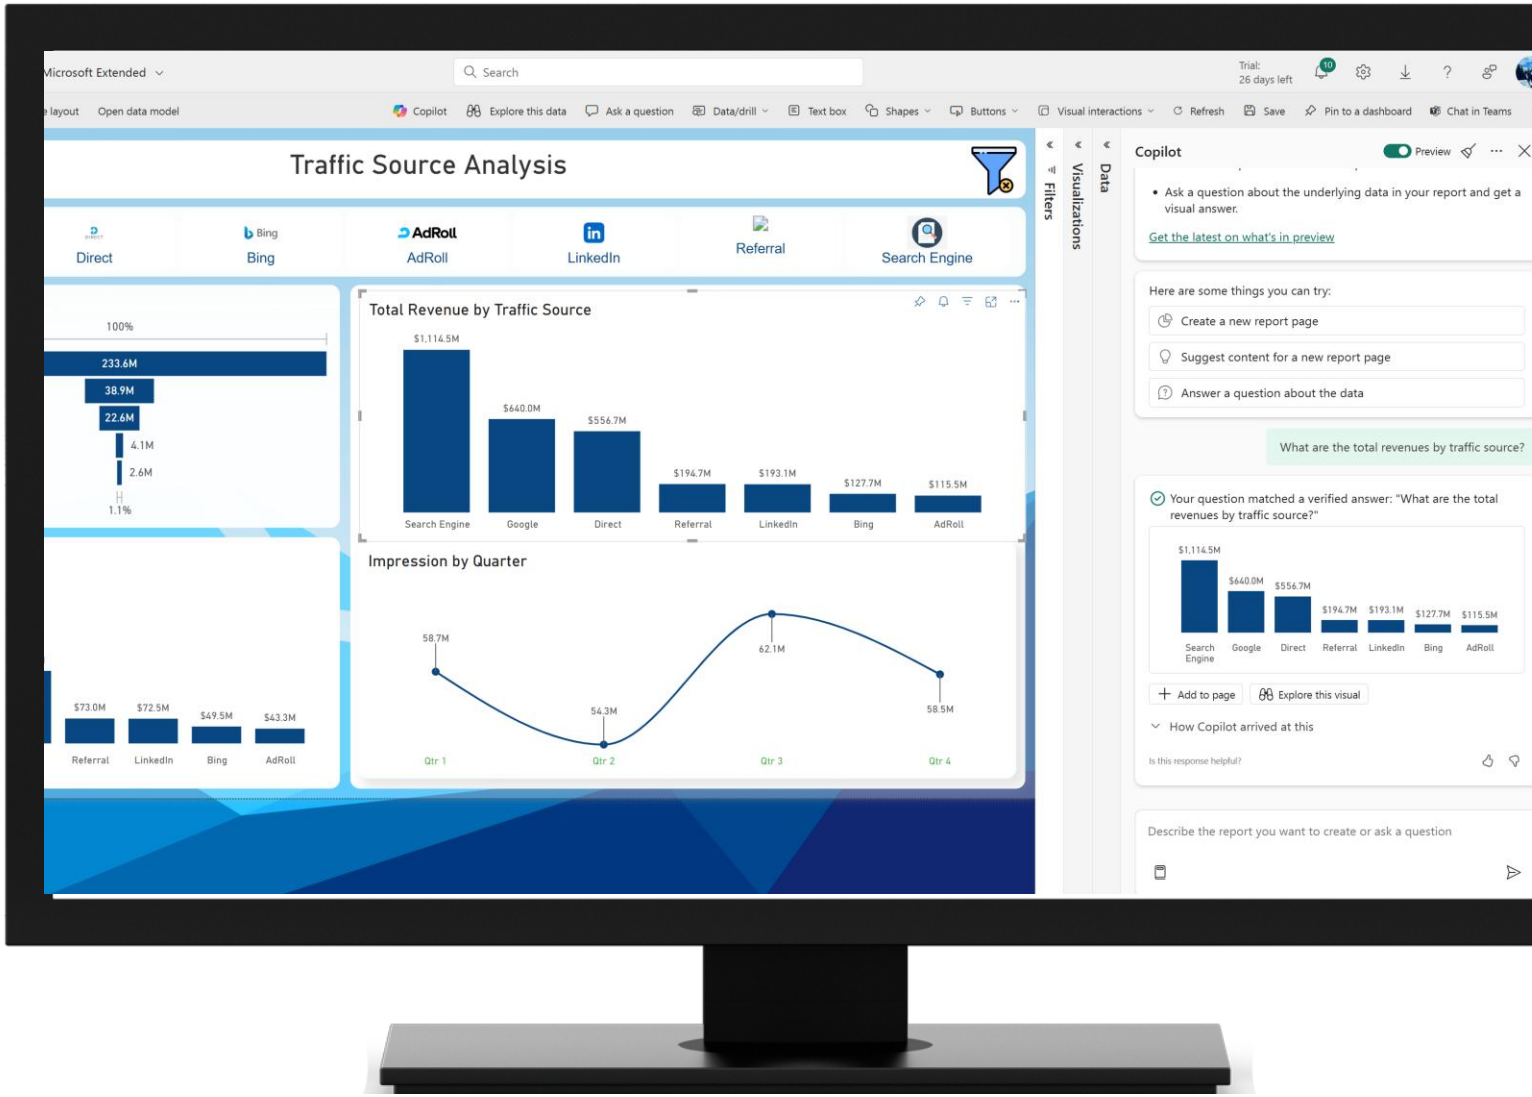


Benefits

Verified Answers

Benefits

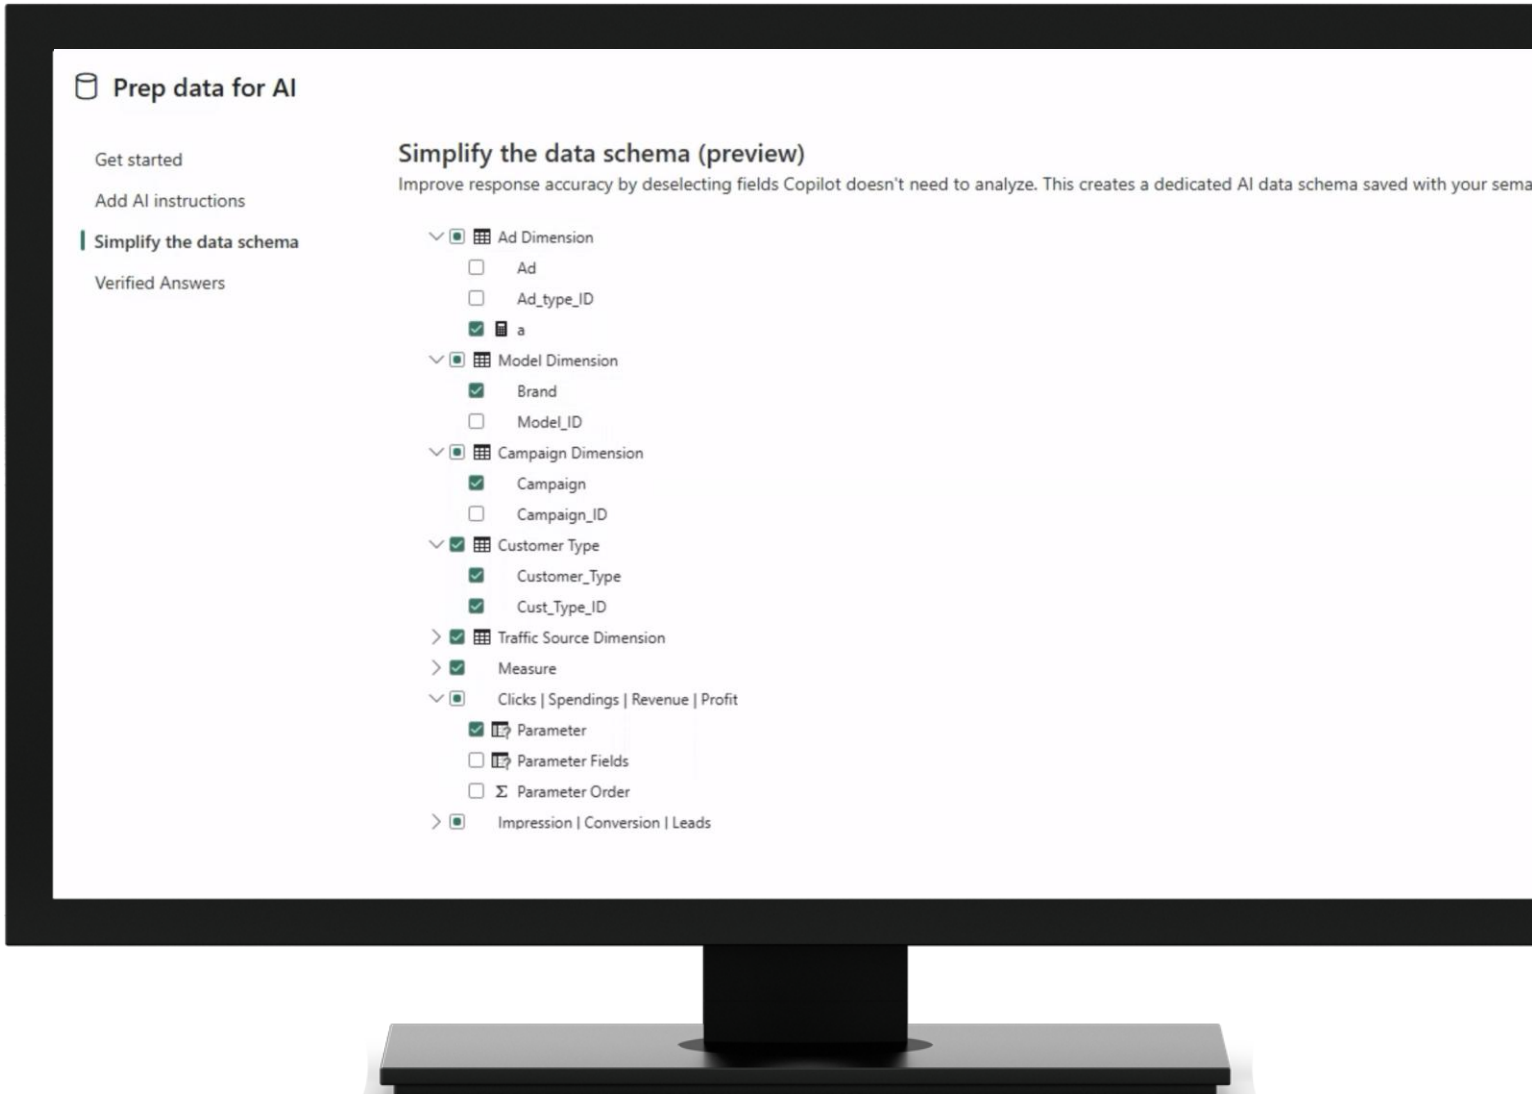
- Accuracy + consistency
- Capture business context
- Help Copilot get smarter over time
- Reuse accurate responses across models



Power BI Copilot Grounding: Schema Selection

Schema Selection

Schema selection allows users to prioritize specific data schemas for more relevant and focused insights.

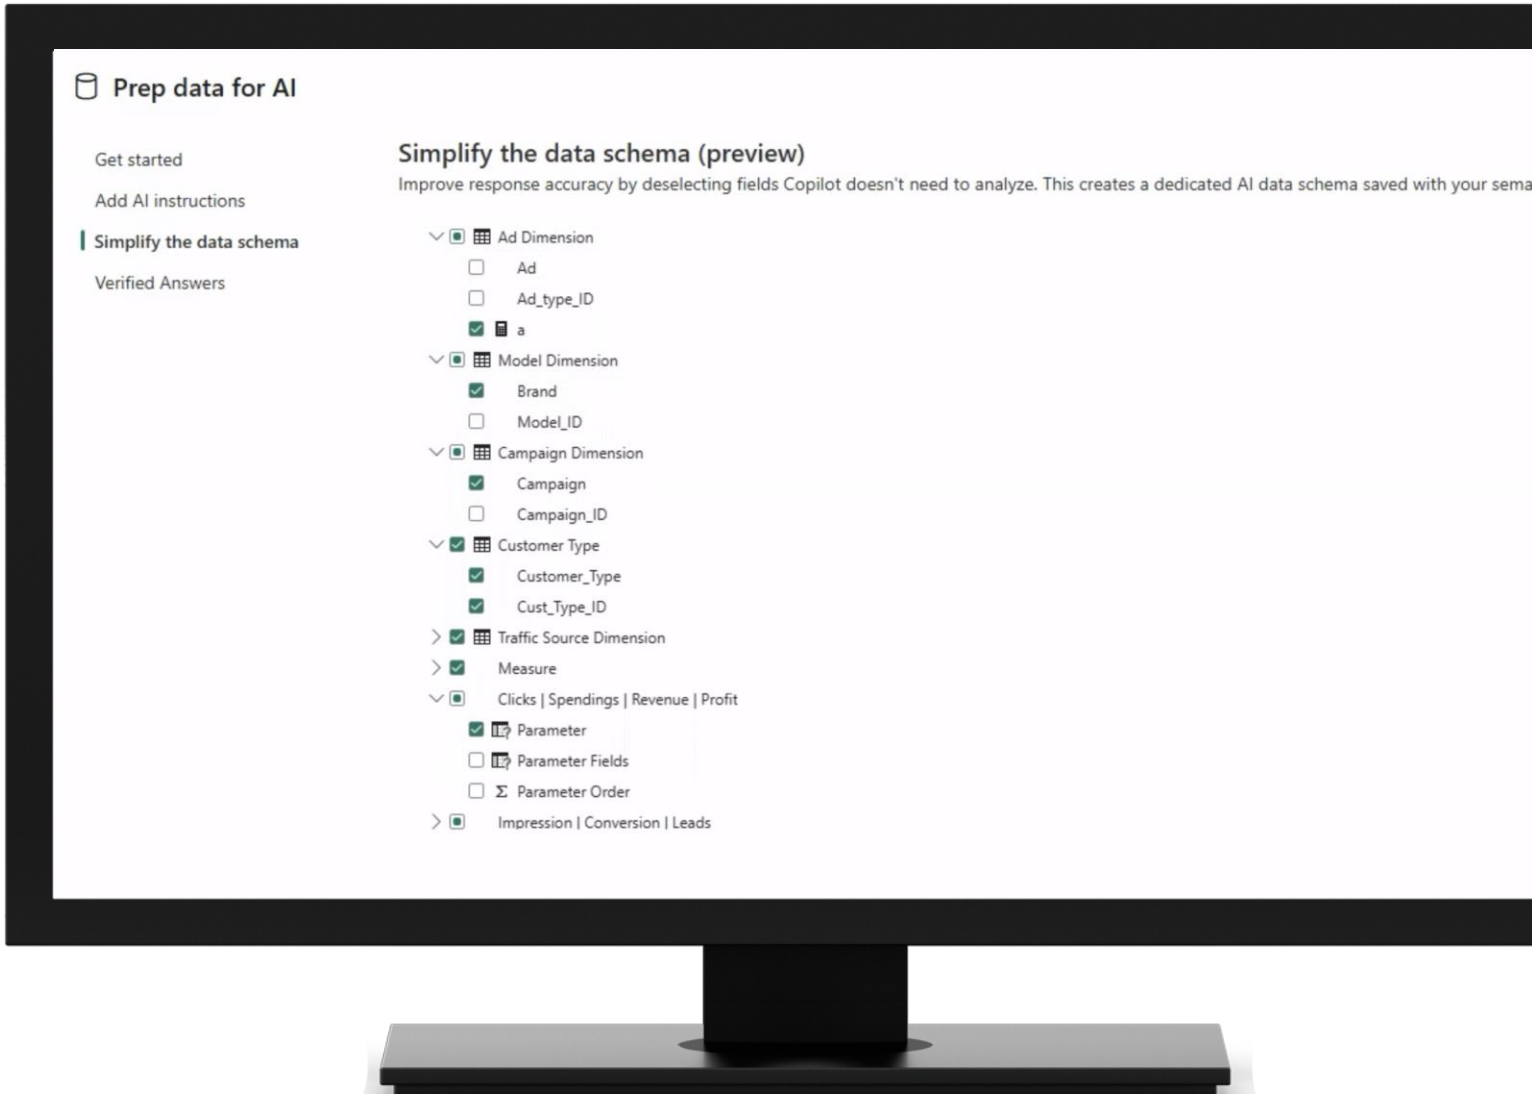


Benefits

Schema Selection

Benefits

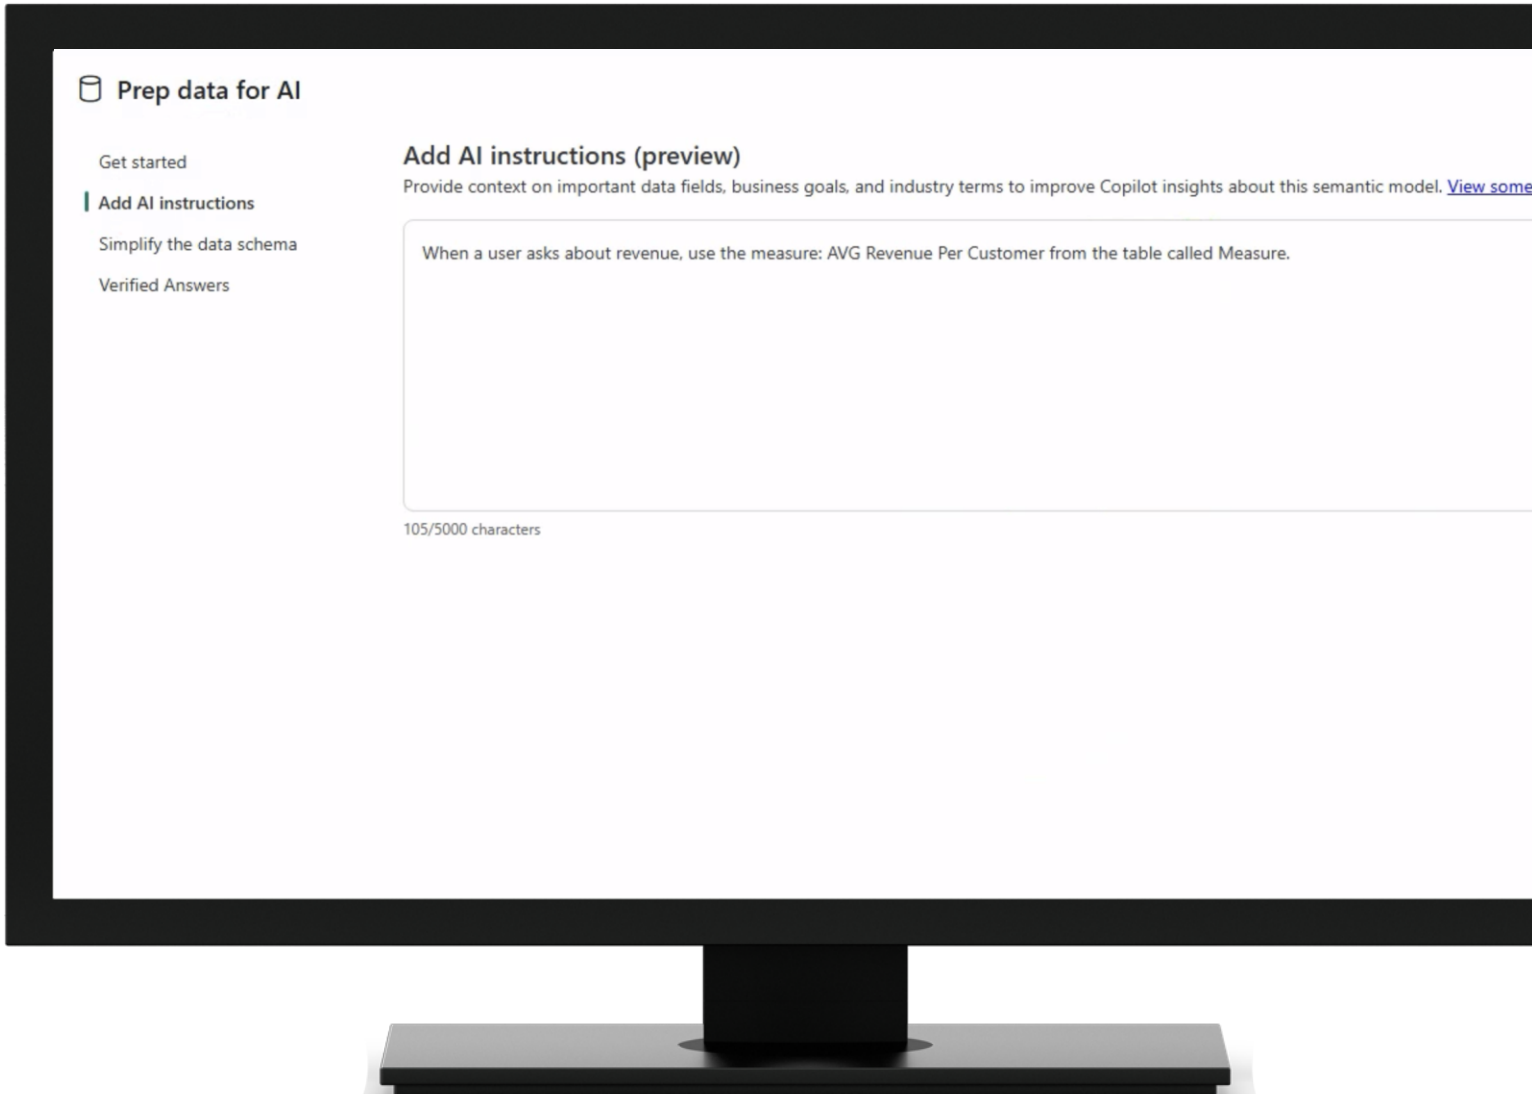
- Column and **measure** level control
- Enhance relevance of specific measures
- Ensure Copilot model responses are friendly and pointing to the right data



Power BI Copilot Grounding: Custom Instructions

Custom instructions

Customize Copilot with instructions for how to best respond to your data.

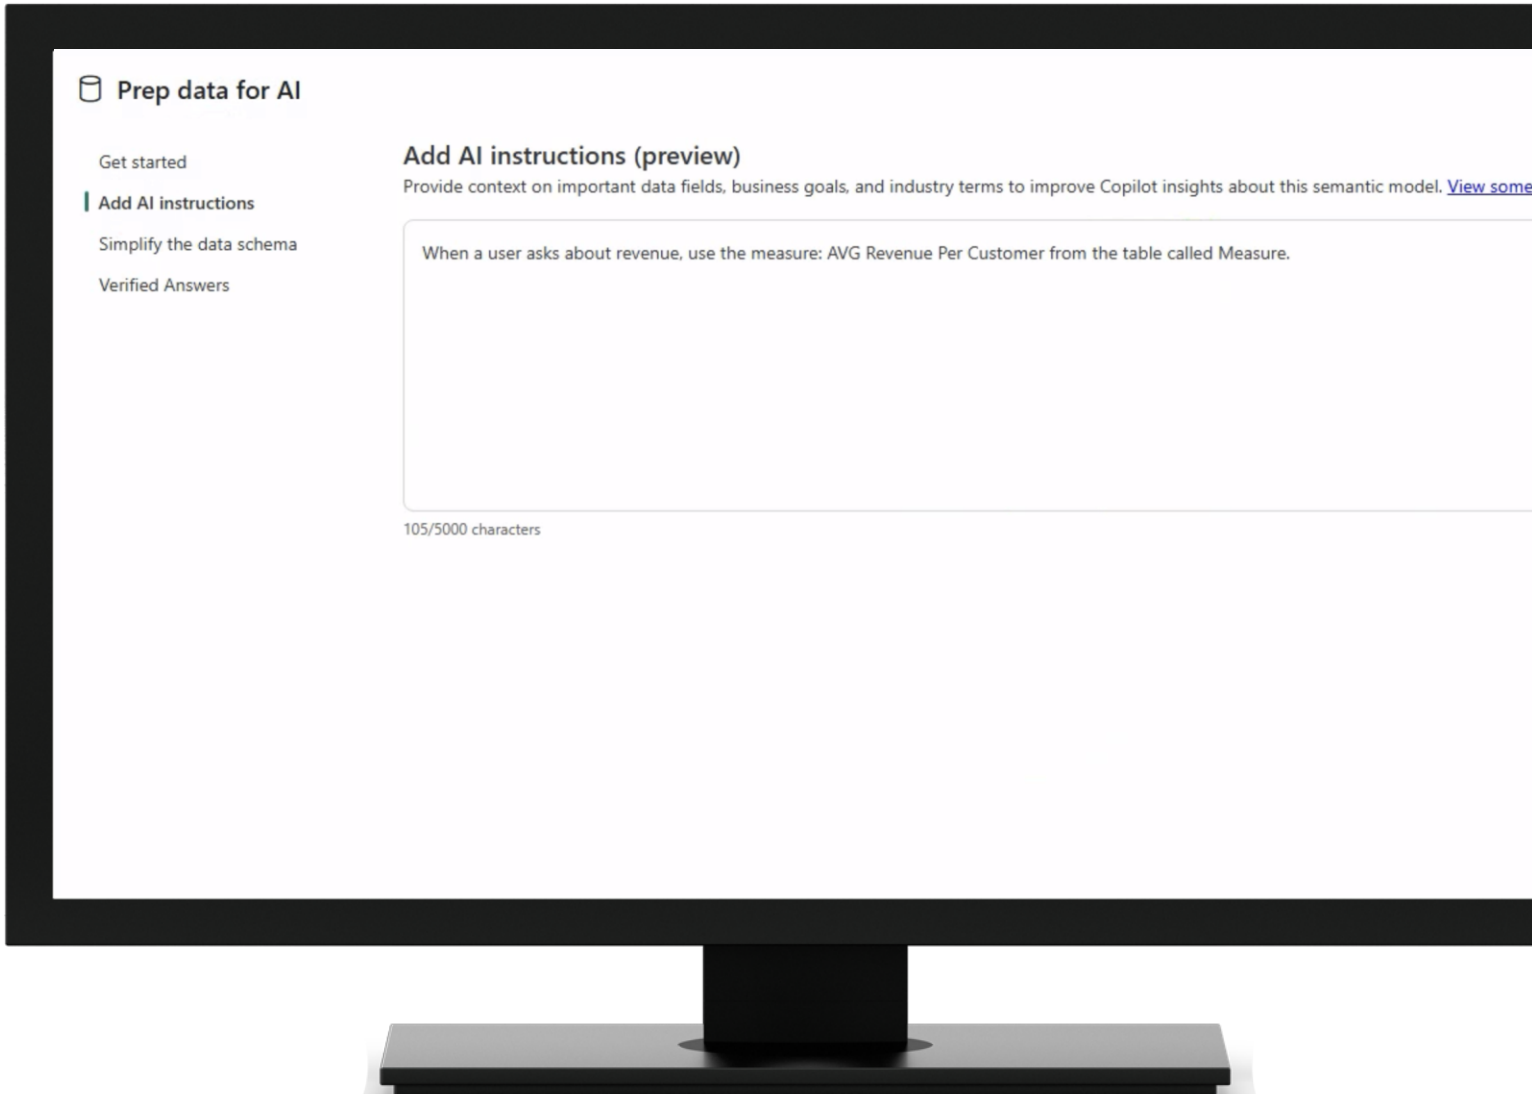


Benefits

Custom Instructions

Benefits

- Personalized Experience: Tailor Copilot's responses with instructions for how to best respond to your data.
- Enhanced Efficiency: Streamline workflows by providing specific instructions that align with user goals.
- Improved Accuracy: Ensure Copilot understands and executes tasks correctly by offering clear guidance.



Best practice and prepare semantic models for Copilot



Best practices – how to get the most of Copilot

- Try and be verbose as it will better understand your request (examples: be specific on field names)
- Steer Copilot, add clarifications to guide Copilot to the correct answer
- Ask follow up questions
- Its an LLM, it can still use ungrounded reasoning
- If something is offensive, it will then return an error as Copilot
- Use thumbs up / thumbs down for feedback to Microsoft

Modeling best practices

Copilot's performance is tightly linked to the quality of your semantic model.

Key practices include

- **Clear human-readable naming conventions**: Avoid ambiguous or overly similar field names. Consistency improves Copilot's ability to interpret prompts accurately.
- **Defined relationships**: Ensure relationships are well-structured and logical (star schema!!). Poorly defined joins or missing keys degrade Copilot's output.
- **Create core measures**: Build essential DAX measures (e.g., Total Sales, Profit Margin, Customer Count) so Copilot has strong building blocks to answer questions.
- **Metadata enrichment**: Add descriptions to model objects, fields, etc. This helps Copilot understand context and improves prompt responses.

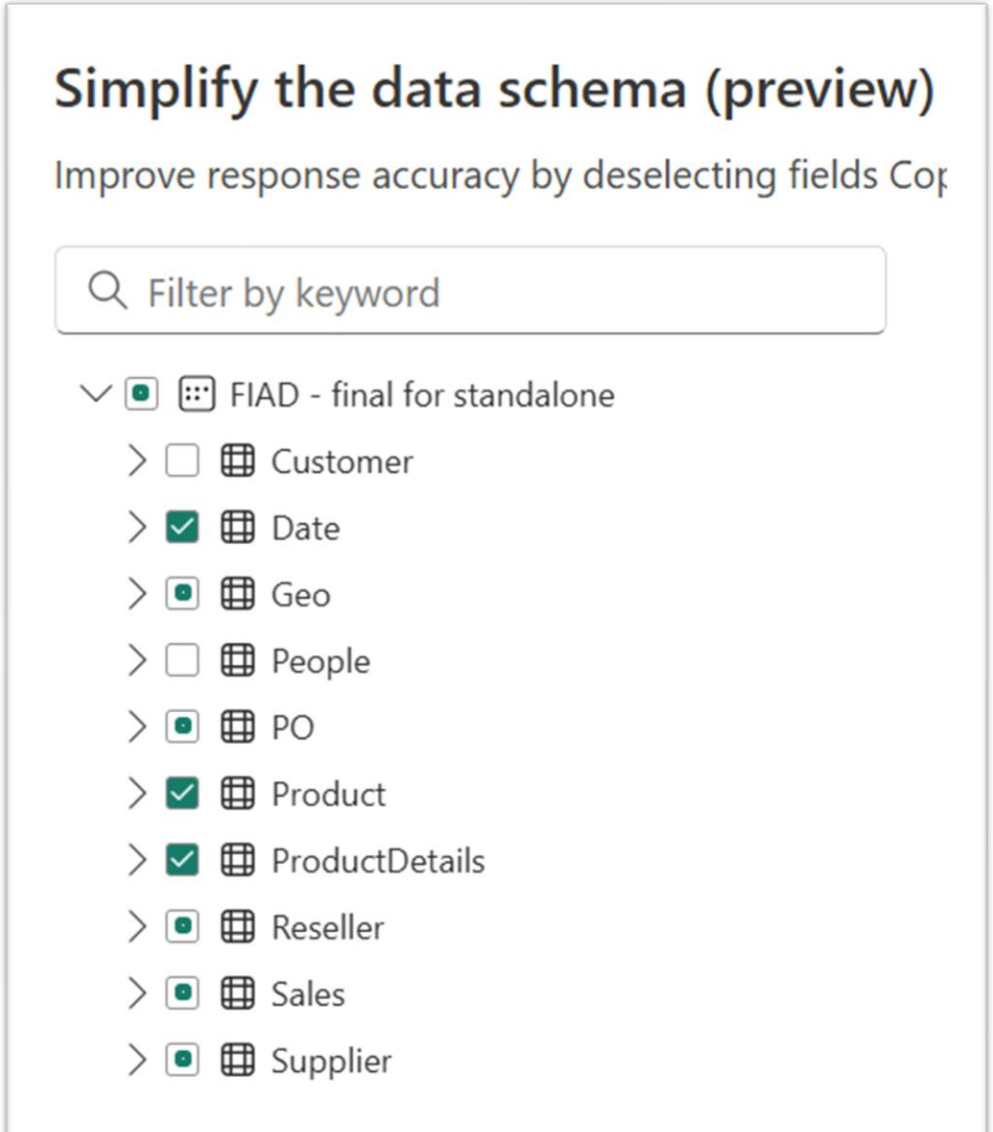
Prep data for AI – Best Practices

- **Simplify the data schema** – 7 tips
- **Verified answers** – 12 tips
- **AI instructions** – 7 tips

Simplify the data schema

Helps Copilot to focus on the right data by only exposing what Copilot needs to see

1. **Visibility level** - Consumers of your semantic model can't see the AI data schema set on the model.
2. **Scope of capability** - The AI data schema isn't considered for Copilot capabilities that don't use the model schema. For example, when Copilot is giving a summary of the report or answering a question directly from the report page.
3. **Hidden state** - All fields available in the semantic model appear in the AI data schema selection pane regardless of visibility set on the semantic model.



Simplify the data schema (continued)

Helps Copilot to focus on the right data by only exposing what Copilot needs to see

4. **Hierarchies always respected** - If a hierarchy is included in the AI data schema, but a column of the hierarchy isn't selected as part of the table, the column is still included as part of the Copilot schema.
5. **Relationships always respected** - Relationships are still respected regardless of the AI data schema set. For example, if two fields are related, and one of the fields is included in the AI data schema, Copilot can still return answers that require that relationship.
6. **No end-user control** - End-users can't disable the AI data schema.
7. **Not a security feature** - Do not use AI data schema as a security feature, it’s a suggestion to Copilot for the schema to use when answering questions

Verified answers

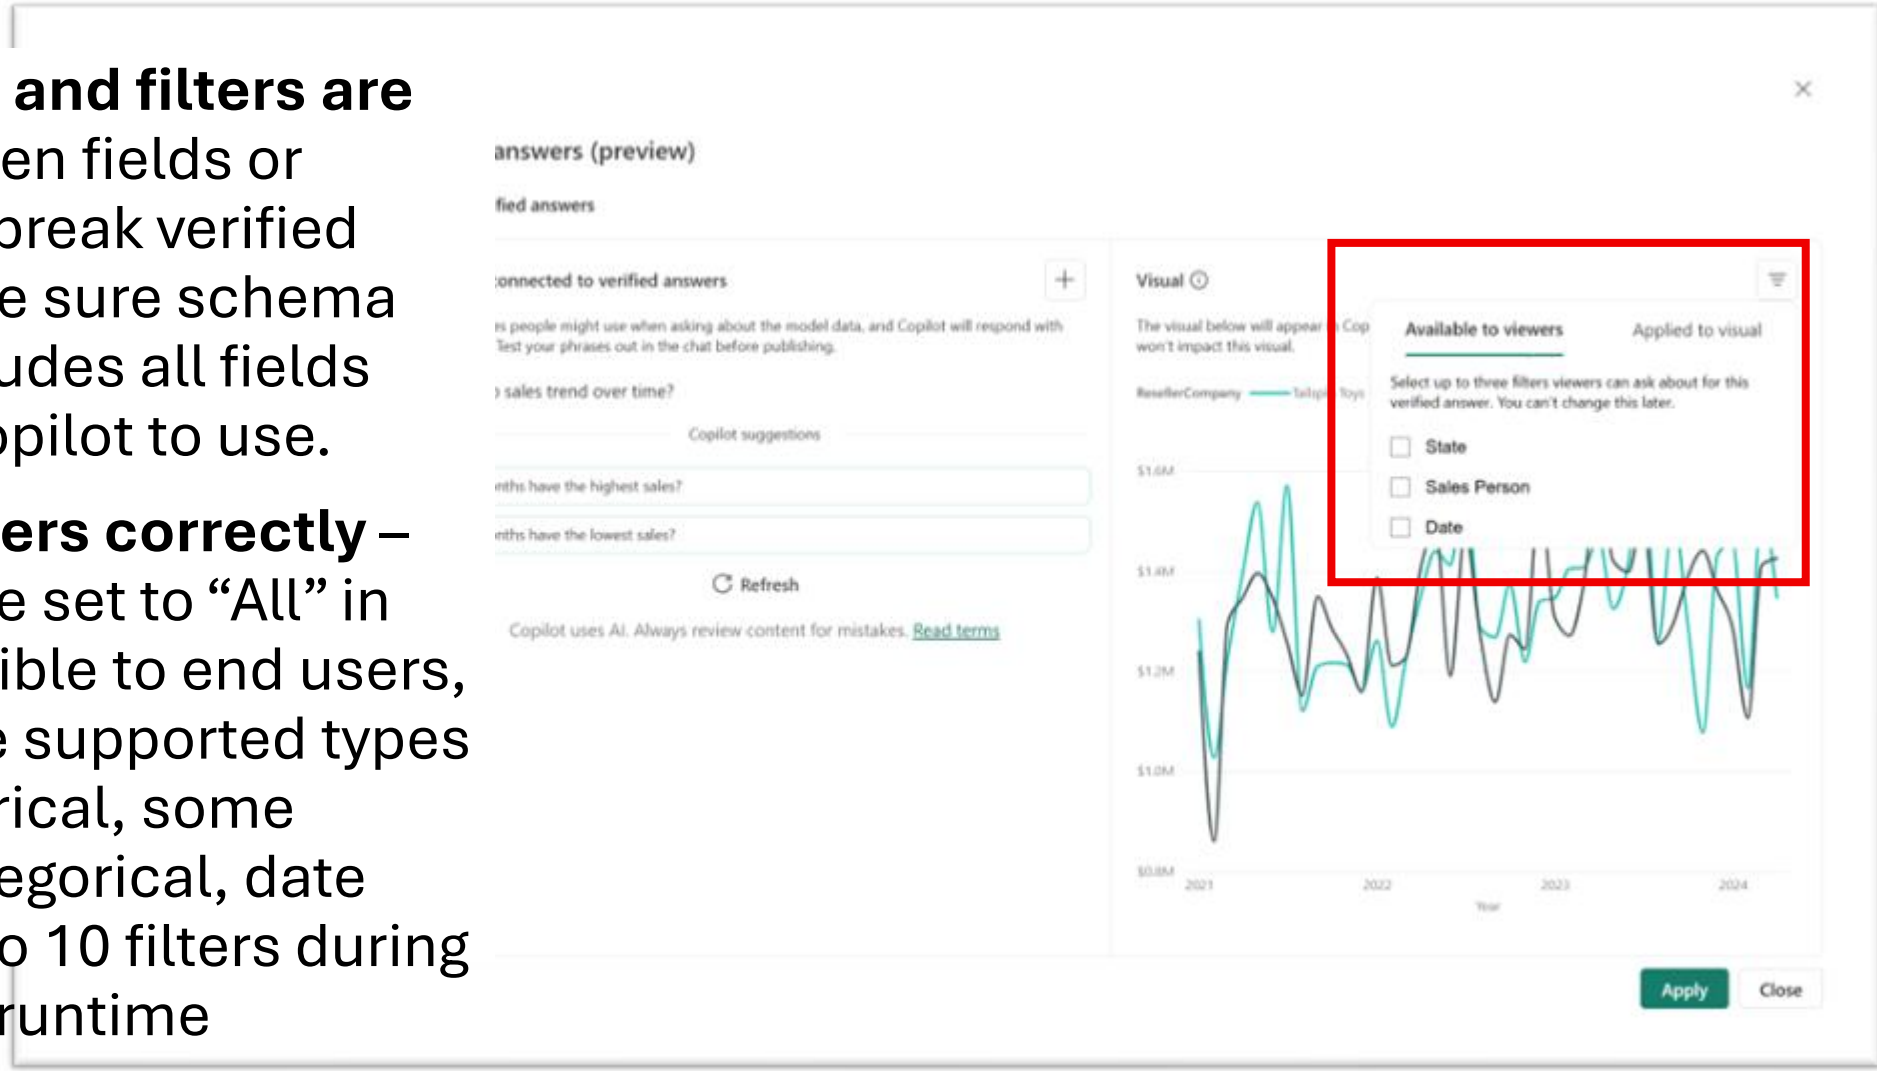
Ground Copilot and have Copilot provide answers defined by the model author

1. **Use both exact and semantic trigger phrases** – Provide 5–7 representative triggers per verified answer to cover different but valid ways users may ask the same question.
2. **Keep semantic matches realistic** – Synonyms, word order changes, and filters in prompts work, but adding/removing fields or swapping measures does not.
3. **Select supported visuals and models** – Verified answers only work with supported visual types (no AI visuals, decomposition tree, etc.) and supported models (Import, Direct Query, Composite, Direct Lake in service).

Verified answers (continued)

Ground Copilot and have Copilot provide answers defined by the model author

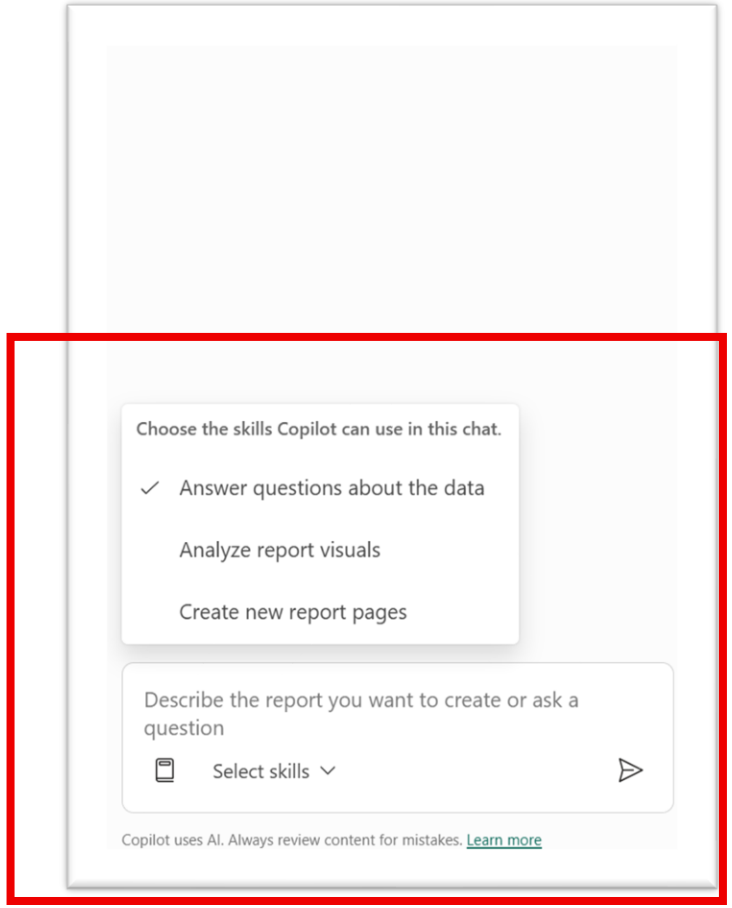
4. **Ensure fields and filters are visible** – Hidden fields or locked filters break verified answers. Make sure schema selection includes all fields you expect Copilot to use.
5. **Configure filters correctly** – Filters must be set to “All” in the report, visible to end users, and within the supported types (basic categorical, some advanced categorical, date range). Limit to 10 filters during creation, 3 at runtime



Verified answers (continued)

Ground Copilot and have Copilot provide answers defined by the model author

6. **Understand answer persistence** – Verified answers are tied to the semantic model, not the original report visual. Editing or deleting the report visual doesn't update the verified answer—manage updates in the Verified Answer setup dialog.
7. **Test carefully with the skill picker** – In desktop, disable authoring skills like "create page" to mimic end-user Copilot behavior and validate your verified answers trigger correctly.
8. **Communicate clearly with users** – Verified answers show a checkmark, matched phrase, and "how Copilot arrived at this." Make sure these elements are accurate and trustworthy for confidence building.



Verified answers (continued)

Ground Copilot and have Copilot provide answers defined by the model author

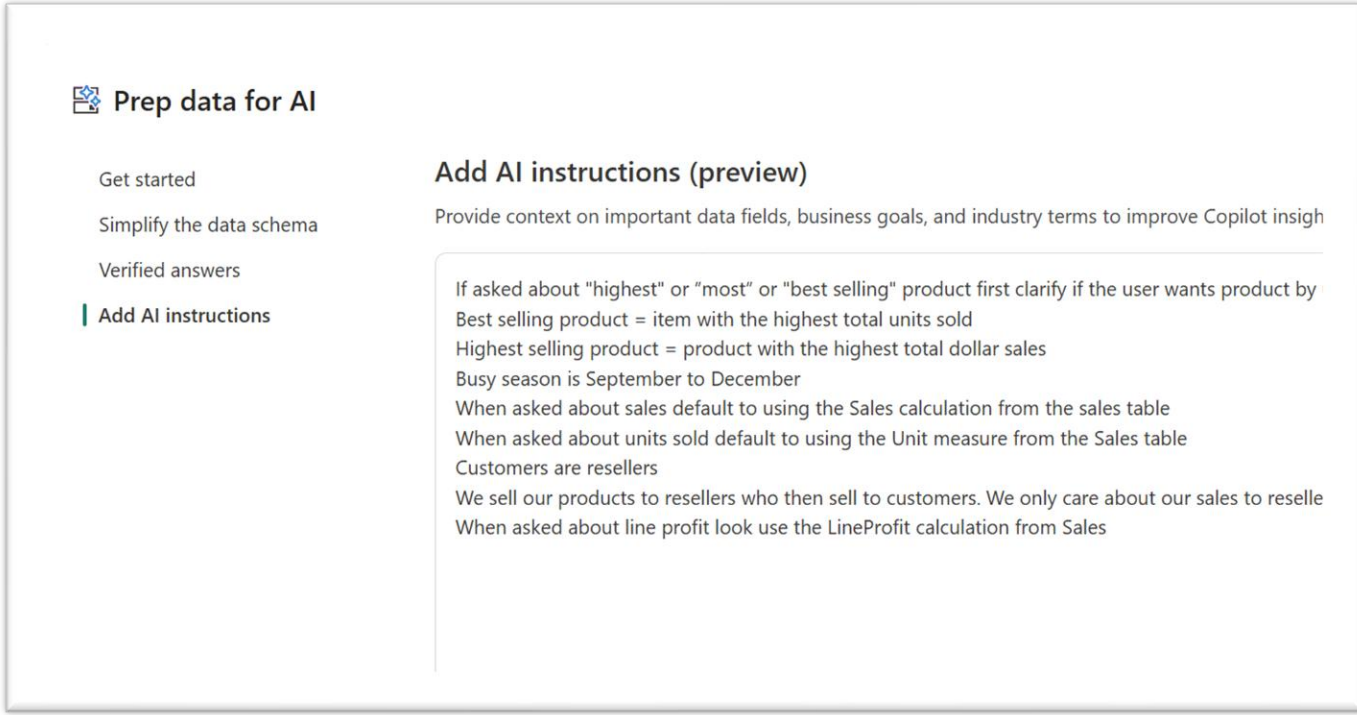
9. **Stay within limits** – Each model supports up to 250 verified answers, 15 trigger phrases per answer, and 500 characters per trigger. Respect these caps for reliability.
10. **Publish and validate in service** – After creating verified answers in desktop, publish to service and allow ~15 minutes for them to appear. A small edit in service (like tweaking a trigger phrase) can speed up sync.
11. **Select the visual first** – the “Set up Verified Answer” option only appears when a visual is selected.
12. **Manage through the Verified Answers dialog** – edits to report visuals won’t update Verified Answers; use the management dialog or recreate them

[Prepare Your Data for AI - AI Instructions]

AI instructions

Give business context and instructions for Copilot to use to provide better answers

1. **Be explicit and specific** – Assume Copilot has no understanding of how to use the data model or business context for your data.
2. **Use analogies and descriptive language** to help the model understand your desired outcome.



AI instructions (continued)

Give business context and instructions for Copilot to use to provide better answers

3. **Avoid ambiguity** – If something should be emphasized - or avoided - be clear. The more details and supporting content you can provide Copilot, the better.
   Example: For Total Active Partners, use the measure "Monthly Active Partner Count" (do NOT filter on "Customers" table).
4. **Group related instructions** – Organizing instructions by theme or purpose (that is, date logic, key metrics, industry terms) can help Copilot better understand your intent. Including structural elements, such as sections, hierarchies, and headers, can be very beneficial.
5. **The order in which you author instructions can impact the output received** – Testing out different variations, orders, examples, and wording can help you get the correct outputs

AI instructions (continued)

Give business context and instructions for Copilot to use to provide better answers

6. **Break down complex instructions into simpler steps** to improve clarity and reduce errors.
7. **Keep instructions focused** – Depending on the model and instructions set, sometimes less is more. Conflicts and complexity in instructions can cause confusion for the large language model (LLM).

When I prepare my data for Copilot, which capabilities are affected?

Capability	AI data schemas	Verified answers	AI instructions	Descriptions
Get a summary of my report	No	No	**Yes**	No
Ask a question about the visuals on my report	No	**Yes**	**Yes**	No
Ask a question about my semantic model	**Yes**	**Yes**	**Yes**	No
Create a report page	No	No	**Yes**	No
Search	No	**Yes**	No	**Yes**
DAX query	No	No	**Yes**	**Yes**

FAQ and Troubleshooting guides for Preparing data for AI:
Frequently Asked Questions about Preparing Data for AI - Power BI

Configuring and grounding data for Power BI Copilot: **Prep data for AI (preview) – available now… And what is next**

Verified Answers

Ensure **quality** and **consistency** of responses for the most critical business topics.

Schema Selection

Limit the scope of data that Copilot can answer questions about

Custom Instructions

Provide business context and help Copilot guide how to talk about your data

You empower your AI agents with the same knowledge and context



CONTEXT ENGINEERING FOR AI AGENTS

Context lives with the item

1

Stored with the item

Context sits on the Fabric item itself, not in a separate tool.

PROVIDE ITEM SPECIFIC CONTEXT

2

Follows its lifecycle

It moves, versions and retires with the item — always in sync.

AUTHORED ONCE, REUSED EVERYWHERE

3

Consumed by all agents

Enrichment turns it into agent-ready context, reused everywhere

TAILORED FOR AGENTS

Why it matters: Context that lives with the item gives better answers out of the box for that item — and it becomes common knowledge, shared with everyone who works with it.

Where we're starting:

Topics

Supply a variety of 'topics' which are dynamically selected based of relevance to the question

---

Remove ambiguity and provide clear business guidance for specific topics, with explainable reasoning steps

---

Consistent across agents – the same guidance is used everywhere the source is used

Templates

Reusable query templates - Verified, parameterized queries stored with the source, so every agent inherits proven patterns

---

Faster answers, under 15 seconds, results return in seconds instead of waiting on the agent to reason out queries

---

Consistent across agents – the same guidance is used everywhere the source is used

CONTEXT ENGINEERING FOR AI AGENTS

Setting semantic model context

1

Work with Copilot

Identify the type of context to author and allow Copilot to provide suggestions

2

Iterate and approve

Review and iterate on the format, details and shape of context

3

Save and consume

Save context to the semantic model

File Edit Selection View Go Run Terminal Help

FabricPM (Workspace)

Update

Problems Output Debug Console Terminal

Terminal

Create Power BI Report Topic - GitHub Copilot

Current Sessions Issues Pull requests Gists

> use the context authoring skill to help build topics for this report: 14:54
> https://daily.powerbi.com/groups/4efdf7e0-e438-4d92-968a-a112593d3fde/reports/ae5b2918-a66a-4409-a21a-9d97a5e313dd/ReportSection?experience=power-bi

> *Thought for 1s*

○ **Asking user** Waiting for response

Copilot needs information.
I resolved **Campaign Analysis (Prepped)** to the **Campaign Analysis** semantic model and confirmed authoring access without writing anything. The report contains five analysis pages, existing AI instructions and verified answers, but no existing topics.

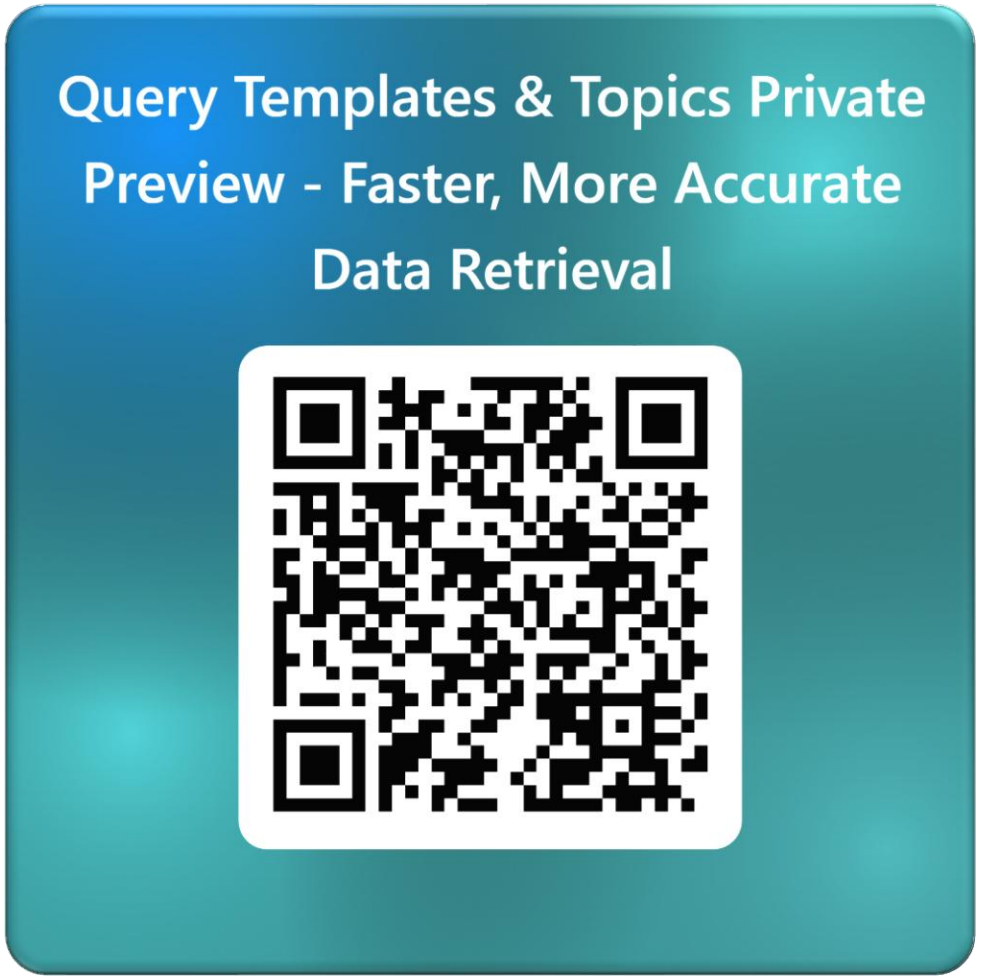
I found three focused opportunities that would improve future answers. Choose one or more to develop:

Recommended topics:
Ranked by expected value. Selecting a topic begins clarification and drafting; nothing is saved without explicit approval.
> ✓ **1. Lead Source vs. Traffic Source — prevent acquisition analysis from using the wrong source concept**
- 2. Spend taxonomy — distinguish Market Spend, Cost, and Total Spendings
- 3. Funnel-stage interpretation — clarify impressions, clicks, visits, leads, and conversions
- Other (type your answer)

↑/↓ select · **space** toggle · **enter** accept · **ctrl+d** decline · **esc** cancel

master* 0 0

Context Authoring: Private Preview



Supports:

Authoring via fabric skill in CLI

Support for semantic models, Data Agents, and semantic models + lakehouses within a Data Agent

---

Consumable via Copilot chat, Power BI Copilot, or within a Data Agent

Rate This session

Are your semantic models ready for Copilot?

