VOLUME I · FIRST EDITION · 2026

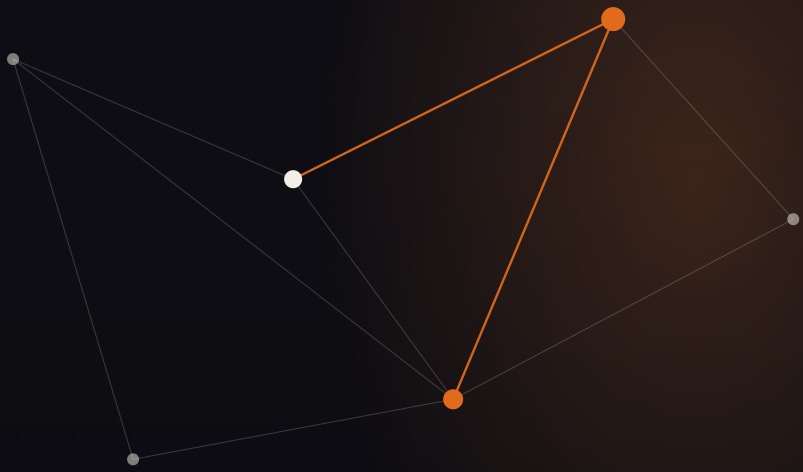

From Reports to *Reasoning*.

The Business Analyst's guide to Agentic AI — and the quiet end of the dashboard era.

POWER BI · EXCEL · MICROSOFT FABRIC

Syed Hussnain
Tahir Sherazi

VOL. 01 / PRACTITIONER SERIES

FROM R → R 2026

A PRACTITIONER'S FIELD GUIDE

From Reports to *Reasoning*.

For every analyst who stayed late on a Friday to fix a broken refresh.

VOLUME ONE

From Reports to *Reasoning*

The Business Analyst's Guide to Agentic AI

POWER BI · EXCEL · MICROSOFT FABRIC

Syed Hussnain Tahir Sherazi

First Edition · 2026

FROM REPORTS TO REASONING
The Business Analyst's Guide to Agentic AI — Volume I.
First edition published in 2026.

DISCLAIMER
This book describes products and features as they were known in early 2026. Software changes quickly. Always check the latest official vendor documentation before making business decisions. The examples in this book are for learning purposes only.

CONTACT
contact@syedhussnain.com · syedhussnain.com

PRINTED IN 2026 · PRACTITIONER SERIES · VOL. 01 · ISBN PENDING

Contents

VOL. I · 10 CHAPTERS

FRONT MATTER

Preface.

I have spent years building reports. I have stared at long Excel sheets at midnight, cleaned messy data, fixed broken formulas, and tried to make Power BI dashboards that nobody opens. If you have done this work, you know the feeling. You spend most of your time preparing data. You spend a small part of your time finding insights. And the business still asks for more.

Then a new wave came. People started talking about "agents." Not robots, but small AI helpers that can think, plan, and act. At first, I thought it was just hype. After all, we have heard big promises before. Big data. Self-service BI. Augmented analytics. Each gave us some value, but none removed the heavy lifting.

Agentic AI feels different. When I first asked an agent to read my data and write a summary, it did it in seconds. When I asked it to find the reason behind a sales drop, it did not just show me a chart. It opened tables, ran queries, and gave me a written answer with evidence. It made mistakes too — and I had to check its work. But the change in speed was real.

This short book is my attempt to share what I have learned. It is written for people like me — analysts, BI developers, finance staff, and team leads who use Power BI, Excel, and Microsoft Fabric every day. I will not pretend that dashboards will disappear tomorrow. But the way we build, share, and act on reports is changing fast. If we do not change with it, our work will become slower and less useful than what an AI agent can do in minutes.

My goal is simple. I want you to finish this book and feel ready. Ready to ask better questions. Ready to use Copilot and Fabric Data Agents with care. Ready to design your next report so an agent can use it well. And ready to lead your team into the next chapter of business intelligence.

I have kept the language plain on purpose. There is no need for heavy academic words when we are talking about real work. If a sentence is too long, I cut it. If a term is technical, I explain it before I use it. Read at your own pace, and try the small examples on your own laptop.

A short note on what this book is not. It is not a Power BI manual. It is not a deep technical reference on prompting or LLM internals. It is a working analyst's map of a fast-moving territory — written by someone who has had to ship reports under deadline pressure, not someone watching from the sidelines.

Where I make claims about Microsoft features, dates, or behaviour, I am writing from what is generally available or in preview at the time of writing. Some of these details will shift in the months after publication. The patterns, however, will hold.

> *The best learning happens when you put your hands on the keyboard.*

— S. H. T. Sherazi, 2026

FRONT MATTER

Why This Book.

Now.

In 2024, most BI teams were still building static dashboards. By 2026, the rules have changed. Copilot now sits inside Power BI, Excel, Word, PowerPoint, and Microsoft Fabric. Fabric Data Agents answer questions in plain English over enterprise data. Multi-agent orchestration in Copilot Studio went generally available in April 2026. The legacy Power BI Q&A feature, based on older natural-language search, is being retired in December 2026. None of this is a small update — it is a shift in how reports are made and used.

Why does this matter for you? Because the skills that made you valuable five years ago are not the same skills that will make you valuable five years from now. Cleaning a column in Excel, writing a DAX measure, or formatting a chart will become tasks that an agent can do faster than any human. The valuable work will move to design, strategy, and judgment.

WHAT YOU WILL LEARN

→ How conventional reporting tools break under modern questions.
→ What "agentic AI" really means, in plain words.
→ How Power BI, Excel, and Fabric Data Agents fit together.
→ How to build a small agent workflow at work.
→ How to handle hallucinations, leaks, and bad outputs.
→ How to grow your career when agents do the heavy work.

How to read this book

You can read it from start to finish, or you can jump to the chapter that matters to you most. Each chapter ends with a short *Try this* section — tasks you can do in your own Microsoft 365 or Fabric tenant. If you do not have access, the free Copilot and Fabric trial cover most ideas.

I will use "agent" and "AI agent" to mean the same thing. I will use "Copilot" when I talk about Microsoft's product brand. When I say "conventional reporting tools" I mean the tools we have used for the last 15 years: Excel pivot tables, SQL reports, and dashboards.

Let us start with the most important question: *why are these old tools no longer enough?*

CHAPTER ONE

01

The End of the Dashboard Era.

For fifteen years the dashboard was king. It pulled data out of the back office and onto the desk of every team leader. Now its weakest habits — passivity, silence, expense — meet a faster opponent.

HISTORY | PAIN POINTS | A DAY IN THE LIFE | WHAT CHANGES

A short history of the dashboard

Twenty years ago, a "report" was a printed sheet of paper. A clerk would run a query on Friday, paste the numbers into Word, and walk it to the manager's desk on Monday. By Tuesday, the numbers were already old. We laugh about it now, but that was the standard for most of the business world.

Then came the dashboard. Tools like Excel pivot tables, Tableau, QlikView, and later Power BI made it possible to refresh numbers daily, weekly, or even live. Suddenly the manager could see sales, costs, and stock levels on one screen. This was a real step forward.

For about fifteen years, the dashboard was king. Every company had a "single source of truth" project, every CFO had an executive dashboard, and every BI team had a long backlog of report requests. Dashboards became the symbol of a data-driven company.

Why the model is breaking

The truth is, dashboards have always had problems. We just learned to live with them. As data grew bigger and decisions had to be faster, those small problems became big walls.

FIG. 1.1 — THE FOUR FAILURES

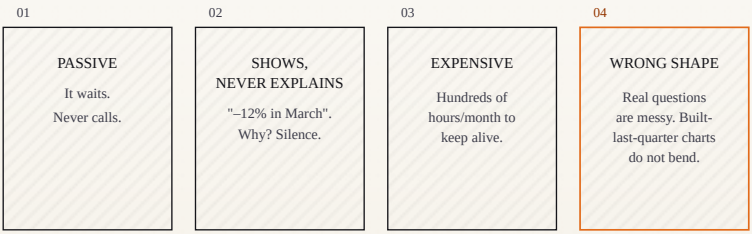

Each pane fails on its own. Together they form the wall that agents climb over.

A morning in the life of a sales VP

Let us picture a Monday for Sara, a sales VP at a mid-sized retailer with 38 stores in three countries.

DASHBOARD MONDAY

8:00 · Opens "Weekly Sales". Looks fine.

8:20 · CFO asks why Pakistan margin fell.

9:30 · BI team needs 3 days for a deep dive.

Result: "I'll get back to you Thursday."

AGENT MONDAY

7:30 · Agent scans data. Margin in PK –4%.

— Sugar cost +18%. Shelf prices not updated.

— Drafts a one-page brief with action.

10:00 · Sara approves the fix from her phone.

> *One waits to be opened. The other works while you sleep.*

What is changing right now

Microsoft, Google, Salesforce, and many smaller vendors all push in the same direction. Reports will not be static pages. They will be living systems that listen, think, and act. The shift began around 2024 with the first Copilot features inside Power BI. By 2026, agentic features are general inside Excel, Word, PowerPoint, and Fabric. The Power BI April 2026 update brought Copilot to mobile. The standalone Power BI agent is in preview. Fabric Data Agents have reached their second year of growth.

What this means for the analyst

Some read these news stories and feel afraid — they worry AI will take their jobs. I don't. In every shift — paper to Excel, Excel to BI, on-prem to cloud — the people who learned the new tool first did better, not worse. The tasks that fade are the ones we already complain about: cleaning columns, copying numbers, rewriting the same DAX measure. The tasks that grow are the ones we enjoy.

In short, the dashboard is not dying because dashboards are bad. It is fading because something better has arrived. Our job is to learn that new tool — and to build the next layer of value on top of it.

FIG. 1.2 — FROM SHOW TO SOLVE

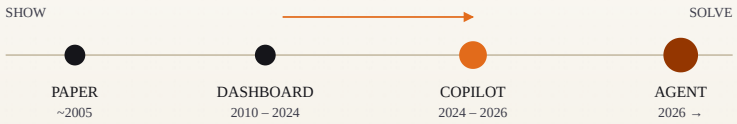

TRY THIS

1. Open the most-used dashboard in your team. Write down three questions it cannot answer.
2. Ask the same questions of Power BI Copilot or Microsoft 365 Copilot. Note the time saved.
3. Share the result with one teammate. The first conversation will surface your team's early adopters.

CHAPTER TWO

02

Understanding Agentic AI in Simple Words.

A chatbot answers. An agent acts. The difference is four small parts working in a loop — brain, memory, tools, and a goal.

BRAIN MEMORY TOOLS GUARDRAILS

From chatbot to agent

Most of us first met AI through a chatbot. You type a question, the bot types an answer. One round, then it stops. A chatbot is like a vending machine — press a button, get a snack. There is no plan, no follow-up, no memory of yesterday.

An agent is different. It has a goal, a set of tools, and the ability to take steps in order. You give it a job — "find out why our top product is losing market share" — and the agent breaks the job into smaller steps. It picks the right tool for each step. It checks the result. It tries again if it fails. When it is done, it tells you the answer and shows its work.

> *A chatbot answers. An agent acts.*

The four parts of an agent

FIG. 2.1 — ANATOMY OF AN AGENT

Put these four parts together and you get an agent that can take a vague human request and turn it into a clear, finished task.

How an agent thinks: a small example

Your boss walks up and says, "*Find out why our top SKU sales dropped in Q1.*" Here is how a Fabric Data Agent handles it.

FIG. 2.2 — A REASONING LOOP

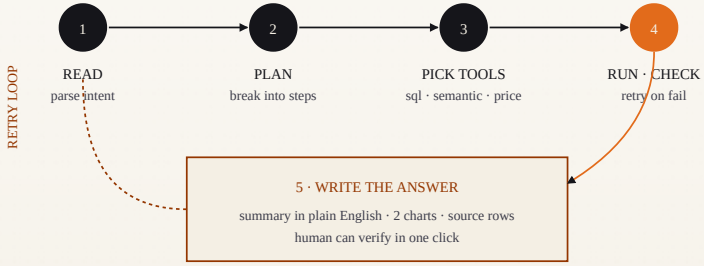

The whole loop takes a few seconds. The same task done by a human might take an afternoon.

Agentic vs generative AI

You will hear both words a lot. They are related but not the same.

- **Generative AI** creates new content from a prompt — a picture, a paragraph, a code snippet. One round and it stops.
- **Agentic AI** uses generative AI as the brain but adds tools, memory, and a loop. It plans, acts, fails, retries.

> *Generative AI gives you a draft. Agentic AI gives you a finished job.*

Why agents matter for reporting

Reports are a perfect job for agents because reports are made of small, repeatable steps: pull data, clean it, model it, calculate KPIs, write a summary. Today's analyst spends most of the day on wrangling. With agents, the ratio flips.

What agents cannot do (yet)

I want to be honest. Agents are not magic. In 2026 they still struggle with several things.

- They can **hallucinate** — the brain writes a fact that is not true. Check key numbers.
- They can **pick the wrong tool**. Similar table names confuse them.
- They can be **slow on big data**. A 10 TB warehouse needs careful prompts and good indexes.
- They can be **insecure**. A poorly designed agent leaks data through prompts.
- They depend on **good metadata**. If tables have no descriptions, the agent guesses.

MENTAL MODEL

Treat an agent like a smart but new team member. Powerful, fast, willing — but in need of context, review, and clear scope. Never the final word on a high-stakes decision.

TRY THIS

1. Open Microsoft 365 Copilot or any free agent. Ask it to summarise your last week of emails.
2. Notice the steps. Did it list folders? Ignore spam? Could you explain how it reached the answer?
3. Try the same prompt twice. The answer may differ. AI output is not deterministic — a key idea for the next chapter.

CHAPTER THREE

03

From Static Reports to Living Decisions.

The report is no longer the end product. The decision — and the action that follows — is.

THREE STAGES · ASK · ACT · APPROVE · AUDITABILITY · MYTHS

The flow we know

Source systems push data into a warehouse. The warehouse feeds a semantic model. The model powers a report. The report is shared on a portal. The user opens the report when they have time. If the report is missing something, they raise a ticket and wait. This flow is linear and slow.

The new flow: from data to action

In an agentic setup, the same data still lives in the warehouse — but the path to the user changes shape. The agent pulls just the part it needs, when it needs it. It does not need a fixed report. It builds the answer on the fly. And it does not stop at the answer. It can send an email, post a Teams message, update a row in Dataverse, or open a ticket. The output of analytics becomes a *doing*, not just a *knowing*.

FIG. 3.1 — TWO FLOWS

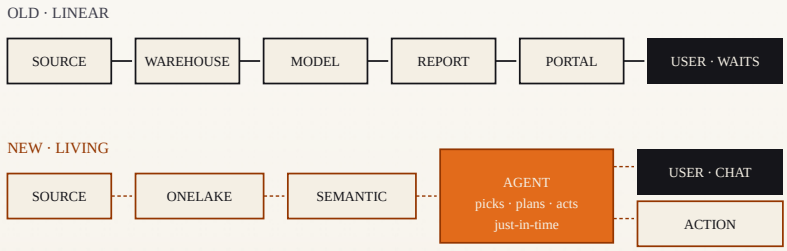

Old: each step takes hours or days. A change at the start breaks things at the end.

New: the agent pulls only what it needs — and finishes the loop with an action humans approve.

"From static reports to living decisions." — Ch. 3

Three stages of agentic adoption

Most companies will move through three stages. Use them to plan your own roadmap.

FIG. 3.2 — THE ADOPTION LADDER

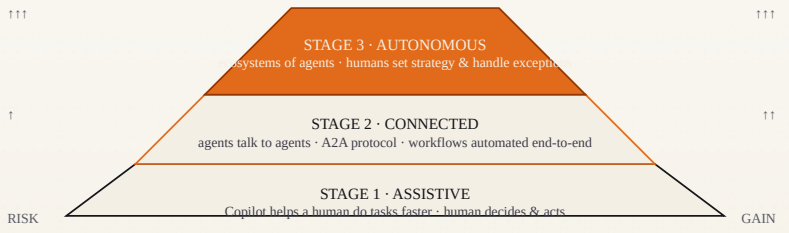

Start at Stage 1. Most teams reach Stage 2 by late 2027. Stage 3 is the 2029–2030 horizon.

What changes for the report itself

In the new world, the report is not a fixed PDF or a single Power BI page. It is a flexible answer made just-in-time for the user.

- **Conversational** — the user asks in plain English.
- **Contextual** — the agent knows who is asking and what they care about.
- **Composable** — it picks the right charts, tables, sources.
- **Actionable** — the user approves a follow-up step inside the same chat.
- **Auditable** — every claim links back to source data and the prompt that produced it.

Auditability is not optional. In a regulated industry, you must show how a number was reached. Modern agentic platforms keep a full log of prompts, tools used, and rows touched. This log is your friend.

How users will work

A common pattern is the *ask-act-approve* loop. The user asks. The agent acts. The user approves — or sends the agent back with a follow-up.

FIG. 3.3 — ASK · ACT · APPROVE

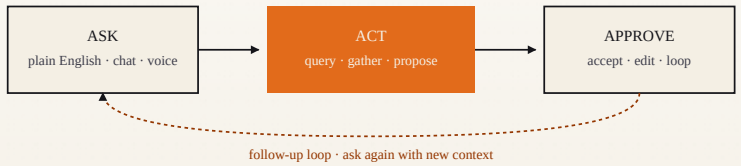

A small case study: the morning brief

A regional bank in the Middle East replaced its daily PDF brief with a Fabric Data Agent. Every morning, branch managers got a Teams message saying, "Good morning, here is your day." It had three parts: top three risks today, top three opportunities, and any urgent compliance items. Each item linked to the source data in OneLake.

Before this, the same brief took two analysts four hours every morning. After the agent went live, the analysts spent that time on deep-dive work the agent could not do. Adoption rose from 35% to 88% in eight weeks.

> *"We did not save money on people. We made our people work on what only people can do." — Head of BI at the bank.*

Common myths to drop

- **Myth 1**: Agents will fully replace BI teams. — They replace tasks, not teams.
- **Myth 2**: A bigger model is always better. — A smaller model with the right tools beats a giant model with bad data.
- **Myth 3**: Once built, an agent runs forever. — Agents drift like data drifts. They need tuning.
- **Myth 4**: Agents work best alone. — They thrive when sharing context with other agents and humans.

Each myth has the same root: a hope that the agent is a finished product, not a living system. Treat your agents the way you treat a junior analyst on their first month — useful, fast, but still in need of supervision, examples, and clear scope. The best teams pair every agent with a named human owner who reviews logs weekly.

TRY THIS

1. Pick one weekly report your team produces by hand.
2. Write down the steps that go into making it.
3. Ask: which of these steps could a Data Agent or Copilot do for us in chat? The list itself will guide your next 12 months.

CHAPTER FOUR

04

Power BI in the Agentic World.

Three agentic surfaces now sit beside the classic report — in-report Copilot, the workspace agent, and the Fabric Data Agent. Each is built for a different user. Together they cover question to answer to action.

COPILOT IN REPORTS | WORKSPACE AGENT | DATA AGENT | MODELING MCP

Where Power BI fits today

Power BI is still the most-used BI tool in the Microsoft world. It has the data model, the visuals, and the sharing layer that millions of users rely on. The tool is not going away. What is changing is the way it is used.

In 2026, Power BI sits inside a bigger family. The semantic model is shared with Microsoft Fabric. Visuals are reachable from Teams, Excel, and Copilot Studio. And on top of all that, three new agentic features sit beside the classic report.

FIG. 4.1 — THREE AGENTIC SURFACES

Copilot inside reports

When you open a Power BI report in 2026, you see a Copilot panel on the right. You click, type a question. Copilot looks at the report's data model, the visuals, and the filters you have applied. It writes an answer based on what it sees.

```
PROMPTS THAT WORK WELL
→ "Summarise this page in three bullets."
→ "Why did revenue drop in March?"
→ "Top five products by margin in the South region."
→ "Suggest a new visual that shows growth by quarter."
```

MOBILE CHAT WITH REPORTS

Since the April 2026 update, the Copilot in Power BI Mobile supports full back-and-forth chat. You stand in a meeting, ask on your phone, get a chart, ask a follow-up, get a deeper view. A real change for field staff who don't sit at a desk all day.

The Power BI agent (preview)

This is more powerful than the in-report Copilot. You don't start from a report — you start from a workspace. The agent searches across many reports, semantic models, and even paginated reports. It picks the right source by itself. From December 2026, it is also the official replacement for the legacy Q&A feature. Plan your migration this year.

Fabric Data Agents

Data Agents live in Microsoft Fabric but can be called from anywhere — Teams, Copilot Studio, a custom app, or a website. A Data Agent has a fixed scope: a set of tables, a set of semantic models, and a list of allowed actions. It is read-only. It cannot change data. It can only answer questions and send the answer to a tool.

```
A TYPICAL SCOPE
Tables — sales fact, customer dim, product dim
Models — Sales semantic model · "Retail-Prod"
Tools — read-only SQL · semantic query · OneLake read
```

The discipline of a scoped agent is what makes it production-ready. Generic chat against your whole tenant is exciting in a demo and dangerous in practice. A Data Agent with three tables and four sample questions can ship in two weeks and run for years.

Treat the scope as a contract. When users ask questions outside it — and they will — the agent should say so plainly and route them somewhere else. Drift in scope is the most common reason an agent that worked fine in month one starts misbehaving in month six.

Choosing the right tool for the job

FIG. 4.2 — DECISION GRID

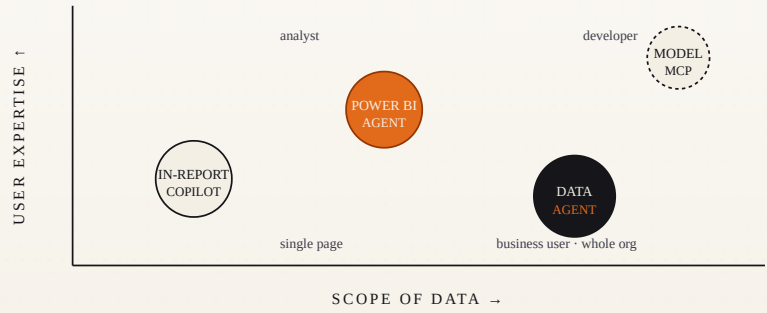

The new monthly review — a practical example

In the old world, the analyst spent two days building the deck. In the new world, prep takes one hour.

1. Monday morning: open the "Monthly Finance" Power BI report.
2. Open the in-report Copilot. Ask: "Summarise this month vs last month, with three risks and three highlights."
3. Copilot writes a clear summary and adds two new visuals to the page.
4. Click "Export to PowerPoint" — Copilot fills the deck with the right slides.
5. Review, edit two slides for tone, save the file.

What used to take 16 hours now takes 60 minutes. The analyst still adds judgment, context, and human framing. That part will not be replaced. The data plumbing is gone.

Tips for getting good answers

Copilot and the Power BI agent are smarter when your data model is clean. A few small habits make a big difference.

DO

- Use clear, full table names ("FactSales" > "F1").
- Add descriptions to every measure.
- Mark the date table with the right setting.
- Keep column names stable.
- Hide internal columns from the report view.

DON'T

- Use cryptic abbreviations or codes.
- Leave measures undocumented.
- Have two tables with similar names.
- Rename columns mid-project.
- Mix internal IDs with display columns.

> *A model is now read by both people and AI. Both want clarity.*

What about other BI tools?

Tableau, Qlik, and Looker have all built agentic features in 2025–2026 — Tableau Pulse, Qlik Answers, Looker Conversational Analytics. Same general idea: a chat interface, approved sources, insights without a fixed dashboard. If your company doesn't use Power BI, the lessons still apply.

TRY THIS

1. Open an existing Power BI report and click the Copilot icon.
2. Ask three questions of different types: a summary, a "why" question, and a "what if."
3. Note where Copilot did well and where it failed. The failures often point to data-model issues worth fixing.

CHAPTER FIVE

05

Excel Reborn — Copilot & Agentic Actions.

Excel never died. It just learned to act. From formula helper to multi-step agent inside the workbook — and a quiet bridge to Fabric, Power BI, and Python.

5 USE CASES | ANALYZE IN EXCEL | PYTHON · GA | VBA → SCRIPTS

Excel is not going anywhere

Some predicted that Excel would die when Power BI arrived. It didn't. Today, Excel is still the most common reporting tool in the world. Almost every analyst, accountant, and operations manager opens Excel before they open anything else. The reason is simple: Excel gives you total control. You can shape data the way you want, in seconds, without asking IT.

What is changing is what Excel can do with help from an agent. Since April 2026, the agentic features inside Excel went generally available. Copilot can now plan and run multi-step actions inside a workbook — bigger than the formula suggestions we saw in 2024.

From formula helper to agent

You give it a goal, and it does the work in steps. You can say: "*I have 5,000 transactions on Sheet1. Group them by region, calculate margin, build a pivot, and add a chart.*" Copilot will read your raw data, detect the columns, suggest a clean structure, build a calculated column for margin, create the pivot, add a chart, and show you what it did — letting you keep, change, or undo each step.

FIG. 5.1 — FIVE AGENTIC ACTIONS IN EXCEL

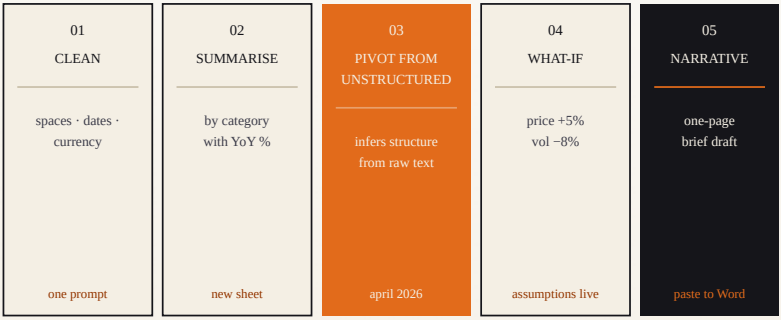

Connecting Excel to Fabric and Power BI

Excel doesn't work in isolation. Through the *Analyze in Excel* feature, you can connect a Power BI semantic model directly to a workbook. Copilot in Excel reads this connection and asks questions of the model in plain English. The same numbers shown in Power BI flow into Excel without copy-paste.

Excel files saved on OneDrive can be picked up by Fabric pipelines. A finance team keeps its working file in Excel, while the rest of the company sees the curated version in Power BI. The two stay in sync. This blurs the old line between "the analyst's Excel file" and "the company's data warehouse."

Excel + Python: a quiet revolution

Python inside Excel went general for everyday users in 2024. By 2026, Copilot can write and run Python code for you, right inside the workbook. Ask "run a regression on these two columns and show the trendline" — Copilot writes the Python, runs it on the cloud, and pastes the result back into a sheet. You don't install anything. For analysts who never learned Python, this is a quiet revolution.

Best practices for Excel + agents

- Use real Excel tables (Insert → Table). Copilot understands tables far better than free-form ranges.
- Name your tables clearly. "tbl_Sales" beats "Table1."
- Keep one fact per cell. Mixing text and numbers in one cell breaks Copilot.
- Avoid merged cells. They confuse the agent.
- Save in OneDrive or SharePoint. Local files limit Copilot's power.

What about old VBA macros?

Many companies still run thousands of macros — some written fifteen years ago by people who have left. Copilot can read VBA, explain what a macro does, suggest a modern replacement, and rewrite it in Office Scripts (the TypeScript-based language for Excel automation). For many small macros, migration is now a 30-minute job, not a 30-day project.

Two steps: ask Copilot to **document** each critical macro — the doc itself is gold. Then plan a slow migration. Replace small macros first. Keep the big ones until you fully understand them.

One word of caution: do not rewrite a macro you cannot test. Old macros often encode rules that were never written down anywhere else. The Copilot rewrite will look cleaner, but it may quietly drop a special case from 2018 that the finance team still depends on. Always run the old macro and the new code on the same input and compare outputs before you switch.

> *The macro that has run quietly for ten years is not bad code. It is institutional memory.*

TRY THIS

1. Take a messy CSV from your work. Open it in Excel.
2. Ask Copilot to clean it, structure it, and add a pivot table by your favourite dimension.
3. Compare the time and effort with how you would have done it manually. This single experience often convinces sceptics.

CHAPTER SIX

06

Microsoft Fabric. The new brain of reporting.

If Power BI is the visual layer and Excel is the personal workbench, Fabric is the brain that holds everything together — one lake, one security model, many agents.

ONELAKE | AGENT ECOSYSTEM | REFERENCE ARCH. | REMOTE MCP

What is Microsoft Fabric, in plain words?

Fabric is a single, all-in-one data platform that brings storage, ingestion, transformation, governance, machine learning, and BI under one roof. Before Fabric, a typical Microsoft data stack used five or six products. Each had its own pricing, portal, access model. Building a pipeline meant gluing tools together.

Fabric replaces this with a unified workspace. The storage layer is **OneLake** — a single, cloud-based data lake for the whole company. On top of OneLake you have a set of "experiences": Data Engineering, Data Warehouse, Data Science, Real-Time Analytics, Power BI, and now Data Agents. They all share the same data, the same security, and the same governance.

Why this matters for agents

Agents need three things: clean data, good metadata, and clear permissions. A unified platform makes all three easier.

- **Clean data.** OneLake stores once, in open Delta Parquet. Agents read without copy-paste.
- **Good metadata.** Descriptions, lineage, and tags flow through Purview. Agents use this context.
- **Clear permissions.** Row- and table-level security pass through to every prompt.

OneLake — the heart of Fabric

Microsoft calls it "OneDrive for data." Every Fabric workspace stores its files in OneLake. Every Power BI semantic model can be backed by OneLake. Every Data Agent reads from OneLake. The big idea is *store once, use many times* — no copying data to a warehouse, then again for ML, then again for the agent. The single copy serves all of them. This saves cost, time, and trouble.

Before OneLake, a typical enterprise had three to five copies of the same fact table living in different systems. Each copy aged differently. Each refresh window was different. Debugging "why are the numbers different in the dashboard and in finance's Excel?" took weeks. With a single OneLake copy, that whole class of problem disappears.

THE FABRIC AGENT TYPES

Data Engineering · helps engineers build & debug pipelines
Data Agents · read-only Q&A over scoped tables
Real-time · watch streams, alert on thresholds
ML Agents · model pick, training, explainability
Copilot Studio · custom, branded, in Teams

A reference architecture for 2026

FIG. 6.1 — A CLEAN AGENTIC FABRIC STACK

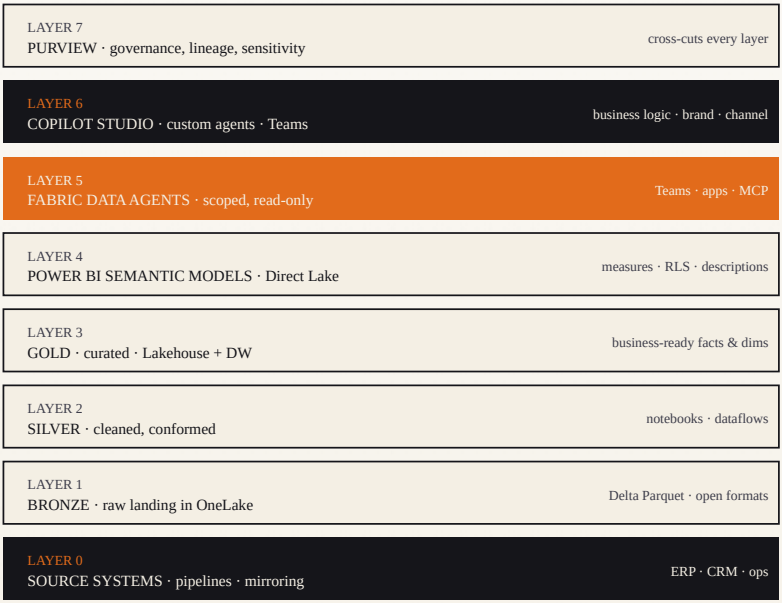

"The cleaner your foundation, the better your agents perform."

Fabric Remote MCP (Preview)

At FabCon 2026, Microsoft announced the Fabric Remote MCP layer in preview. MCP stands for **Model Context Protocol** — the open standard for letting any LLM-driven tool connect to your data and functions in a safe way.

Why does this matter? Today most AI tools only work if they sit inside Microsoft's ecosystem. With Remote MCP, an external tool — say, Claude Desktop or a custom Python agent — can connect to Fabric, ask for data through approved tools, and respect all of your security rules.

> *Fabric Remote MCP is one of the most under-appreciated launches of the year. Pay attention. It will shape how agents and data come together for the next decade.*

Cost and capacity in the agent era

Fabric uses a capacity-based pricing model — you buy F2, F4, F8, and run anything within that limit. Agents don't change the billing model, but they change the usage pattern. A single curious user with an agent can run hundreds of queries per hour, where in the dashboard era they ran one or two.

This is not a reason to be afraid of agents — it is a reason to plan capacity differently. The old planning model assumed predictable, scheduled refreshes. The new one needs to account for bursty, human-driven exploration. Buy a smaller base SKU and lean on auto-scale, rather than over-buying for the peak.

```
GOVERNANCE THAT SAVES MONEY
→ plan for higher peak loads
→ use auto-scaling where possible
→ set per-agent quotas
→ watch the Fabric capacity metrics app daily
→ most 2026 over-spends came from agents running uncapped
queries on large fact tables
```

TRY THIS

1. Sign up for the free Microsoft Fabric trial if you don't have access yet.
2. Create a Lakehouse, upload a small CSV, and build a tiny semantic model.
3. Create a Data Agent on top of that model. Ask three questions in plain English. In 30 minutes you'll see why this is a different way of working.

CHAPTER SEVEN

07

Building Your First Agent Workflow.

Seven concrete steps. Eight weeks. One Sales Insight Agent live in Microsoft Teams. The goal is not perfect — the goal is shipped.

DEFINE | PREPARE | BUILD | GOVERN

Up to now, we have looked at ideas. In this chapter we'll build a small agent workflow from start to finish. Our example: a **Sales Insight Agent** any team member can chat with in Microsoft Teams. It reads a Fabric semantic model and answers questions about regional sales, top products, and risks.

FIG. 7.1 — THE SEVEN-STEP LIFECYCLE

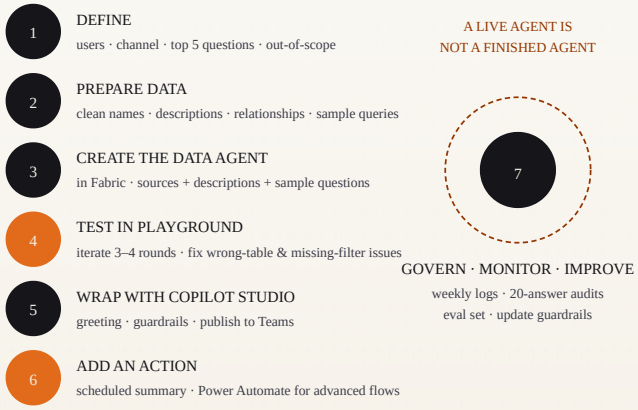

Step 1 — Define the use case

Always start with a clear, narrow use case. Wide goals like "an AI for the whole company" never finish. Narrow goals like "answer five common sales questions in Teams" finish in a few weeks.

```
ONE-PAGE BRIEF

Users · Sales managers and the country head.
Channel · Microsoft Teams.
Top 5 questions · How is region X this month? Top 10
products by margin? Customers at risk of churn? Why did
revenue drop in country Y? What is the Q4 forecast?
Data sources · Sales fact, Customer, Product, Region —
Gold layer in OneLake.
Out of scope · HR data. Finance. Payroll. Anything other
than sales.
```

Step 2 — Prepare the data

No agent works well on bad data. Spend at least half your time here. Pick the four tables. Use clear, full names. Add descriptions to each table and to each measure. Make sure relationships are correct in the semantic model. Run a few sample queries by hand to confirm numbers match your old reports.

> *If a new analyst could understand your model in 10 minutes, an agent can use it well. If your model needs an hour of explanation, the agent will struggle.*

Step 3 — Create the Fabric Data Agent

In your Fabric workspace, create a new Data Agent. The wizard asks for three things: a clear name and description; the four tables from Step 2; sample questions. Behind the scenes, Fabric builds embeddings for your tables, writes hidden prompts that explain your semantic model, and sets up logging and lineage. You don't see this — but it makes the agent reliable.

Treat the description field as if it were the only documentation a stranger would ever read. "Sales fact table at order-line grain, refreshed nightly from the ERP, primary key OrderLineKey" is useful. "Sales data" is not. The agent reads these descriptions to decide which table to query — every word you put there is doing work at runtime.

Step 4 — Test in the playground

Before publishing, test. The Fabric portal has a built-in playground. Type your top five questions one by one and check the answers.

- Wrong table picked when two tables have similar names → clearer descriptions or remove the unused table.
- Wrong total because of a missing filter → improve your semantic model measures.
- "I do not know" too often → more sample questions and better metadata.

Iterate three or four times. Each round takes minutes, not days. Keep a running list of the questions that failed — that list is the start of your evaluation set, which you will use forever after to catch regressions.

Step 5 — Wrap with Copilot Studio

A Data Agent on its own is great, but most users don't open Fabric. They live in Teams. So you wrap the Data Agent inside a Copilot Studio agent that lives in Teams. Open Copilot Studio. Create a custom agent named "Sales Insight Bot." Connect it to your Fabric Data Agent. Set a friendly greeting. Add safety guardrails — block answers about HR, payroll, or anything outside scope. Publish to Teams. Users find it under Apps and add it to any chat or channel.

Step 6 — Add an action (optional but powerful)

Up to now the agent only answers questions. To make it really useful, give it an action. "*Send a one-page summary to my team channel every Monday at 8 AM.*" In Copilot Studio you build an autonomous action that wakes on a schedule, asks the Data Agent for the latest summary, formats it, and posts it. No code needed for simple cases. For advanced ones, use Power Automate.

Step 7 — Govern, monitor, improve

- **Monitor logs.** Fabric and Copilot Studio keep full chat history. Look at it weekly.
- **Track quality.** Pick a sample of 20 answers per week. Score them good / partial / wrong.
- **Improve metadata.** Most "wrong" answers come from unclear column names or weak descriptions. Fix the source, not the agent.
- **Update guardrails.** As people try new prompts, add new rules for new edge cases.
- **Run an eval set.** Keep 30–50 fixed test questions. Run them after every change.

FIG. 7.2 — REALISTIC TIMELINE

TRY THIS

1. Pick the smallest, dullest report you have — the one nobody wants to maintain.
2. Try to replace it with a Fabric Data Agent in two weeks.
3. Even if it doesn't fully work, the lessons set you up for the next, bigger project.

CHAPTER EIGHT

08

Prompt Design, Governance & the New Mindset.

Building the agent is the easy part. Getting good answers from it, day after day, is the hard part. Skip this chapter and the agent will lose user trust within months.

CASE PATTERN | 5 PILLARS | ROLES | MINDSET

Prompt design for analysts

A prompt is just the words you send to an agent. A good prompt is clear, specific, and gives context. A bad prompt is short, vague, and leaves room for guessing.

FIG. 8.1 — THE C·A·S·E PATTERN

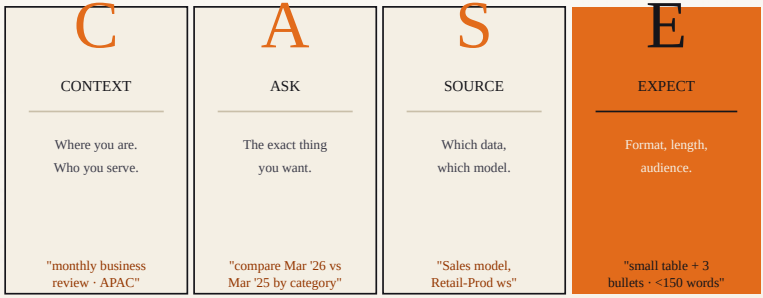

COMMON PROMPT MISTAKES

- Asking many things at once. Break them into separate prompts.
- Using internal codes the agent has never seen. Spell them out the first time.
- Giving no time frame. The agent will guess, often wrongly.
- Forgetting the audience. "For my CEO" or "for a junior analyst" changes the tone.
- Skipping the format. If you don't say "table" or "bullets," you get a wall of text.

Governance for agentic BI

Governance is the set of rules and processes that keep your agents safe, fair, and useful. In an agentic world, governance is the difference between a tool that helps and a tool that hurts.

FIG. 8.2 — THE FIVE PILLARS

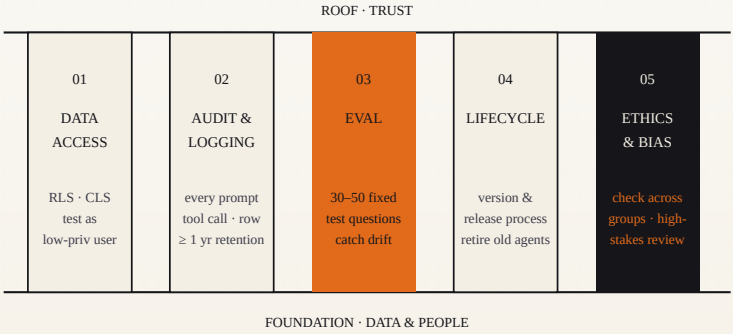

Together, the five pillars carry the roof — and the roof is *trust*. Lose one pillar and the structure leans. Lose two and it collapses, usually publicly.

Roles in the new BI team

↗ GROWING

- Agent designer — picks use cases, writes prompts.
- Semantic model engineer — builds clean models agents can use.
- Governance lead — owns access, audit, risk.
- Evaluation analyst — runs tests, scores outputs.
- Insight communicator — turns agent output into stories.

↘ FADING

- Pure dashboard builder.
- Routine SQL writer.
- Manual data cleaner.

"Fading" doesn't mean "gone tomorrow." It means less needed each year. Use the next two or three years to grow into a new role.

Skills to invest in

1. **Data modelling.** Star schemas, granularity, slow-changing dimensions.
2. **Prompt & conversation design.** A core analyst skill now, not a fringe one.
3. **Storytelling with data.** Agents produce numbers; humans turn them into stories.
4. **Domain knowledge.** The more you know about your industry, the more value you add on top of generic AI.
5. **Light coding.** Python, Office Scripts, and a bit of TypeScript.

A new mindset

- **Trust but verify.** Treat every agent answer as a draft. Check the key numbers before you act.
- **Curiosity over fear.** Try the new tools. Break them. Learn what they can and cannot do.
- **Teach others.** The fastest way to build expertise is to explain it.

The analysts who will thrive over the next five years are not the ones who learn the most prompts. They are the ones who keep their hands on the actual data — who still know what a star schema looks like, who can read a query plan, who can spot when an agent's confident-sounding answer is quietly wrong. The agent is a tireless junior. You are the editor.

> *An agent answers in seconds. A trustworthy answer still needs a human who has earned the right to sign off on it.*

TRY THIS

1. Take a recent prompt you sent Copilot. Rewrite it using the CASE pattern.
2. Compare the two answers side by side.
3. Save the better one in a "good prompts" file. Build your own library over time.

CHAPTER NINE

09

Risks, Limits & the Hard Truths.

Every honest book on AI must have a chapter about what can go wrong. Six risks every BI team should plan for — with a one-page checklist.

HALLUCINATION | BIAS | COST | SKILL LOSS

Risk 1 — Wrong answers (hallucination)

The most famous AI risk is when the model makes things up. In Microsoft's BI stack, this risk is smaller — Copilot mostly works on top of approved data, not from open-internet knowledge. But it can still happen. The agent may use the wrong table, apply the wrong filter, or pick the wrong measure when two are similar.

How to reduce it: clean metadata, fewer measures with clearer names, more sample questions, and a strong eval set. How to live with it: train users to always look at the data behind the answer, not just the answer itself.

Risk 2 — Bias hidden in data

If your historical data is biased — say, more sales effort in some regions — the agent will repeat that bias in its insights. It will tell you the same regions are most important, even though that may be a result of past investment, not of the actual market. Bias is not a Microsoft problem. It is a data problem. But agents make it louder, because they speak with confidence.

Risk 3 — Over-reliance on the agent

When something is fast and easy, people use it without thinking. A user who would never have built a report by hand now asks for ten "quick" answers a day, then forwards screenshots to senior leaders. Some of those answers will be wrong, or right but misleading. The first instinct of any senior leader should be to ask: "*How do you know this is correct?*"

A FOLLOW-UP CULTURE

→ every chart asks where the data came from
→ which agent answered, against which model
→ who reviewed it before it left the team

Risk 4 — Cost surprises

Agentic BI can be cheaper than the old way (less ETL, fewer custom reports), but the cost curve is also harder to predict. A single curious user can run thousands of agent calls in a week. If you don't watch capacity carefully in the first months of any rollout, you will get a bill nobody planned for.

Three healthy habits: set up a monthly capacity review, define per-agent quotas, and tell users when the system is in trial mode and when it's not.

Risk 5 — Privacy and over-sharing

The biggest source of leaks in 2025 was not malicious. It was over-permissive access. Sites were opened to "everyone in the company." Files were uploaded with sensitive data. Copilot saw everything its caller had access to, and shared it on request. The fix is not to block AI. The fix is to clean up access rights first.

Risk 6 — Skill loss in junior analysts

If junior analysts grow up only with Copilot and never write a DAX measure or design a star schema by hand, they may never understand what the agent is doing. When the agent gets it wrong, they cannot tell. Companies that take training seriously will keep a healthy junior pipeline; those who don't will face a skills cliff in five years.

The hard truths in one page

Every shift in this industry came with a few uncomfortable facts that teams preferred not to talk about until they were forced to. This one is no different. Here are seven of them, written plainly so you can plan around them rather than be surprised by them.

```
SEVEN HARD TRUTHS
1. Some of your favourite reports will not survive this decade.
2. Some of your favourite roles will change shape.
3. Some of your favourite tools will be retired by Microsoft.
4. Some new tools will fail and be retired too. Don't bet the farm on a single preview.
5. Your data quality issues are now visible to the whole company through Copilot.
6. Your governance gaps are now exploitable.
7. The companies that win will be honest about points 1–6 and act on them now.
```

What the next chapter won't say

You won't read here that "AI will solve all problems." It won't. You won't read that "AI is just hype." It isn't. The truth sits in the middle — and the middle is exactly where useful, careful, well-run BI teams have always lived.

CHAPTER TEN

10

The Road Ahead — 2027 and Beyond.

An honest look forward. What's likely. What's possible. What's wishful thinking. And what should you do this quarter.

2026 → 2030 | LIKELY | POSSIBLE | WISHFUL

Likely · Possible · Wishful

FIG. 10.1 — A HONEST FORECAST

LIKELY (2026 – 2028)	POSSIBLE (2027 – 2029)	WISHFUL (don't bet on)
→ Every BI tool has chat	→ Static dashboards rare	× "100% autonomy" in finance
→ Q&A retires Dec '26	→ Agents draft reviews end-to-end	× Agents that "fix" bad models alone
→ Multi-agent in Teams	→ "Insight engineer" becomes a real title	× Zero-governance, "just trust the AI"
→ Semantic models become a job spec	→ Cross-vendor MCP marketplaces	× One model for every use case
→ Fabric Remote MCP goes GA	→ Smaller, fine-tuned in-tenant models	× The end of analysts
→ AI literacy programs scale across orgs		

Plan against the LIKELY column. Pilot against POSSIBLE. Politely ignore WISHFUL.

The next 90 days

The next 90 days are more important than the next ten years for most teams. Why? Because in 90 days, you can pick a use case, ship an agent, learn the real lessons, and decide if AI is worth the bigger bet. Without that 90-day experience, you'll spend the next ten years debating.

FIG. 10.2 — A REALISTIC 90-DAY PLAN

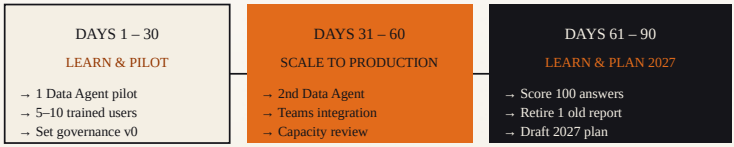

Skip the conferences. Skip the consultants for one quarter. Go build something.

A closing thought

> *Reports always served one purpose — helping people make better decisions. That purpose has not changed in fifty years. The shape of the artefact will keep changing. What stays constant is the need for a thoughtful human in the loop. Be that thoughtful human.*

TRY THIS

1. Block one hour in your calendar this week. Write down: where is my team in this story? What is the next single step we can take?
2. Send it to one colleague. Ask them to do the same.
3. Meet for coffee next week to compare notes. That is how the new BI starts in your company.

Glossary

Forty-two terms every BI professional should know in 2026. Short, plain, useful.

Agent. Software that can plan, decide, and act on its own across multiple steps. In BI, an agent can read data, write a query, render a chart, and send it.

Agentic BI. Business intelligence delivered through agents instead of fixed dashboards.

Audit log. A record of every prompt, tool call, and result, used to investigate problems later.

Bronze, Silver, Gold. Three layers of refinement in a data lakehouse: raw, cleaned, and curated.

CASE pattern. A simple prompt template — Context, Ask, Source, Expect.

Copilot Studio. Microsoft's tool for building custom agents that live in Teams, Outlook, or websites.

DAX. The formula language used in Power BI semantic models.

Data Agent. A Microsoft Fabric agent scoped to a set of tables and models. Read-only by default.

Direct Lake. A Power BI mode that reads OneLake Parquet files directly, no import step needed.

Eval set. A fixed list of test questions used to grade agent quality over time.

Fabric. Microsoft's unified data platform: ingestion, lakehouse, warehouse, ML, BI, all in one.

Fabric Remote MCP. The 2026 preview that lets external tools talk to Fabric using the open MCP standard.

Governance. The rules and processes that keep data and agents safe, fair, and useful.

Guardrail. A rule that blocks unwanted agent behaviour, like answering off-scope questions.

Glossary — *continued*

H through T.

Hallucination. When an agent invents a fact that isn't in the data. Lower risk in BI than in open chat, but real.

Lakehouse. A storage pattern that combines the cheap, flexible nature of a data lake with the structure of a warehouse.

LLM. Large Language Model. The "brain" behind a Copilot or agent.

MCP. Model Context Protocol — an open standard for connecting AI models to tools and data sources.

Metadata. The descriptions, tags, and lineage that tell humans and agents what a piece of data means.

OneLake. The single, organisation-wide data lake at the heart of Fabric.

Power Automate. Microsoft's no-code/low-code workflow tool. Often used to give agents actions like sending emails or updating systems.

Power BI agent. A workspace-wide agent in Power BI, the official replacement for the legacy Q&A feature.

Prompt. The words you send to an agent. Good prompts are clear, specific, and give context.

Purview. Microsoft's governance and compliance product. Tracks data lineage and sensitivity across Fabric, M365, and beyond.

RLS. Row-Level Security. A rule that hides certain rows from certain users.

Semantic model. A friendly layer on top of raw tables — measures, relationships, names — used by Power BI and agents.

Star schema. A modelling pattern with a central fact table and surrounding dimension tables. Still the most agent-friendly shape in 2026.

Tool call. When an agent invokes a defined function — a SQL query, an email send — instead of writing prose.

∴

Author's Final Word

A short letter to the reader.

I wrote this book because too much writing about AI sits at the extremes. One side claims everything will change overnight; the other claims it is all hype. Neither matches the day-to-day work of running a BI team — which is what most of us actually do.

The reality is more interesting. Microsoft has put together, in 2026, the most complete agentic BI stack we have ever seen. Used well, it can save hundreds of hours each month, free analysts for harder work, and bring data closer to the people who need it. Used badly, it can erode trust, leak data, and trick teams into believing numbers that look right but aren't.

The difference between the two outcomes is not the technology. The technology is the same for everyone. The difference is the people: the analysts who learn, the leaders who govern, the teams who keep human judgement at the centre.

If this book helps you and your team have one better conversation this month — about what to keep, what to retire, what to try — it has done its job. Thank you for reading. And please, write to me. I want to hear how it goes.

— Syed Hussnain Tahir Sherazi

Spring 2026

✦

THE AUTHOR

Syed Hussnain Tahir Sherazi

DATA ENGINEER · WRITER ON AGENTIC AI

Syed Hussnain Tahir Sherazi is a Microsoft Certified Business Intelligence Development expert with deep proficiency in Power BI and Microsoft Fabric. He has built his career in environments where mistakes in data are not really a choice. He earned his early stripes in the global smartphone industry, working across brands including Realme, Infinix, Tecno, and Itel, where data driven decision making shaped product launches, user experience research, and after sales performance across Pakistan, Africa, and China. Later he moved to the UK and stepped into the public sector, where data is rough, inherited, and unforgiving, and where reports face directors, elected members, and the general public. He currently works at Leicester City Council, building end to end Power BI solutions, contributing to the migration from on premises SQL Server to Microsoft Fabric, and authoring the council's official Power BI Style Guide. He writes about how analyst work is changing as Copilot, Data Agents, and Fabric reshape the day to day of business intelligence, and about the small daily habits that decide whether an AI rollout becomes a productivity gain or a quiet liability. He holds a Master's in Cloud Computing from the University of Lincoln and a Bachelor's in Software Engineering from the University of Management and Technology. He writes from Leicester.

✉ contact@syedhussnain.com
◉ syedhussnain.com

❦

Colophon

This book was set in *Newsreader* for reading text, with *Inter* for sans-serif passages and *IBM Plex Mono* for diagrams, captions, and running heads. Page diagrams were drawn in plain SVG, printed at 300 dpi on uncoated stock.

First edition · April 2026 · printed in the European Union on FSC-certified paper.

· · ·

£24.99 / €29 / $32
BUSINESS · TECHNOLOGY

Reports are not dying. Reports are becoming conversations. The dashboards of the last thirty years are slowly being replaced by something that listens, reasons, and acts.

— FROM THE INTRODUCTION

For thirty years, business intelligence has meant one thing: build a dashboard, share a link, hope someone looks. In 2026 that whole pattern is being rewritten. Microsoft has put a Copilot inside every Office app, a workspace agent inside Power BI, scoped Data Agents inside Fabric, and an open MCP layer connecting them all.

This is the field guide to what's changed and what hasn't. Inside, you'll find a sober tour of the new stack, a seven-step playbook for shipping your first agent in eight weeks, the CASE pattern for prompts, a five-pillar governance model, and a clear-eyed look at the risks every team must plan for.

- → What replaces the dashboard, and what doesn't.
- → How the Power BI agent retires the legacy Q&A.
- → Fabric, OneLake, and the new Remote MCP layer.
- → A real timeline: pilot in 4 weeks, production in 8.
- → Governance that keeps trust intact.

— A field guide for the next decade of BI —

SYED HUSSNAIN TAHIR SHERAZI
Data engineer and writer on agentic AI in the Microsoft data stack.
Power BI · Excel · Fabric.

SYEDHUSSNAIN.COM
CONTACT@SYEDHUSSNAIN.COM