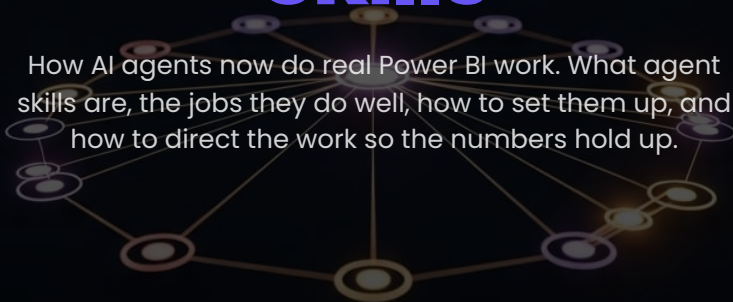


A FIELD GUIDE FROM ENTERPRISE DNA

Power BI Agent Skills

How AI agents now do real Power BI work. What agent skills are, the jobs they do well, how to set them up, and how to direct the work so the numbers hold up.



FRONT MATTER

A note from Sam

If you searched for Power BI agent skills and came away unsure what the phrase meant, you are not alone. It points at three different things at once, and this guide sorts them out.

The bigger story sits underneath the confusion. For years, doing real work in Power BI meant writing the DAX by hand, dragging every visual, and building the model one step at a time. That is changing fast. You now describe the outcome, an agent drafts the measure or the model or the report page, and you verify and approve. The analyst moves from operator to director.

This is real, and it is early. Gartner expects three quarters of analytics content to lean on generative AI for context by 2027. It also expects more than four in ten agentic AI projects to be abandoned by then, most of them on weak governance rather than weak technology. The winners will not be the teams with the most access. Access is becoming universal. The winners will be the ones who adopt it well.

We do not write this from the outside. Enterprise DNA runs its own company on a command center of agents, with a person holding the gate. Every habit in this guide is one we use ourselves. Read it as a field guide. What agent skills are, the jobs they do well, how to set them up, and how to keep the numbers honest.

Sam McKay

Founder, Enterprise DNA

CONTENTS

What is inside

- 1 The shift.** Why Power BI work is changing, and what agent skills actually mean.
- 2 The highest-value skills.** The jobs agents do well, and where a human still signs off.
- 3 How agents reach your model.** The three doors, the Microsoft surface, and your first setup.
- 4 Trust it, then own it.** Verifying agent output, and the analyst's new role.
- 5 Where Enterprise DNA fits.** How we help, and how to start.

PART 1 · THE SHIFT

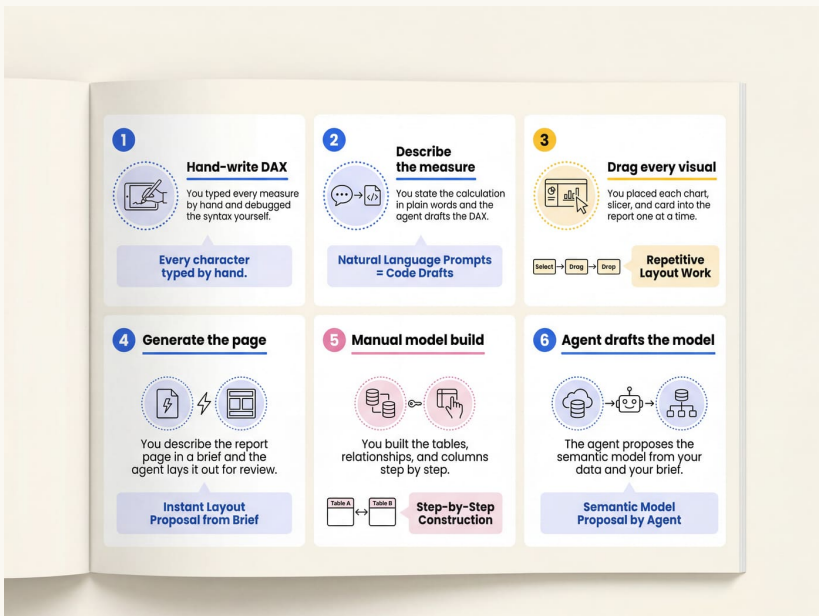
Why this matters now

For years, putting Power BI to work meant doing the work by hand. An analyst wrote the DAX, dragged every visual into place, built the model column by column, and clicked through every step to get from raw data to a finished report.

That is changing. Now the analyst describes the outcome and an agent drafts the measure, the model, or the report page. The analyst reviews the draft, corrects it, and approves. The work still gets checked by a person. The person just no longer types every line.

The analyst stops typing every measure and starts directing the work.

This moves the analyst up a level. The job shifts from operator to director. You spend less time on syntax and layout and more on the question, the logic, and whether the numbers are right. Microsoft framed its Build 2026 announcements as the agentic era of analytics, and the direction is clear.



This is real, and it is early. The tools are in preview, the patterns are still forming, and the good habits are being written right now. The rest of this guide is the practical map. What agent skills are, how to build them, and how to direct the work so the output holds up.

What agent skills actually mean

The phrase agent skill shows up in two places in 2026, from two different companies, and they point at the same idea.

Two definitions, one idea

Anthropic published Agent Skills as an open, cross-platform format in December 2025. A skill is a folder. It holds a SKILL.md file that describes a repeatable capability, plus optional scripts. The agent loads the skill only when the task needs it. The same skill runs in the Claude apps, in Claude Code, and through the Agent SDK.

Microsoft announced Agent Skills for Power BI at Build 2026, now in preview. You prompt an agent to build and refine semantic models and reports. It takes the full path from raw data to semantic model to published report, starting from a plain-language description or even a screenshot.



The unifying idea is simple. An agent skill is a packaged, repeatable capability you hand an agent so it does a specific Power BI job your way, every time. It is a named workflow, not a one-off prompt.

A plain prompt

You ask once and forget it. The next person starts from scratch.

A skill

Reusable, versioned, and shared. The team runs the same job the same way.

PART 2 · THE HIGHEST-VALUE SKILLS

The jobs agents do well

The next four pages walk the Power BI skills where an agent earns its keep. They run roughly in order of payoff, from the DAX you write every day to the audit you keep meaning to run.

Every skill here follows one pattern, so learn it once. **The agent does the production work.** It drafts the measure, edits the model, builds the page, cleans the table. **The analyst verifies the number and owns the definition.** The agent is fast and tireless and occasionally confidently wrong. You are the one who knows what the business means by "active customer" and whether the grand total reconciles.

So read each skill as a division of labour, not a handover. The agent produces a first draft you can inspect. You keep the judgment calls, the metric definitions, and the final sign-off. Nothing reaches a certified dataset without a person approving it.



The agent produces. The analyst verifies and owns the meaning.

DAX and the semantic model

Two skills sit at the top because they touch the core of every report. Writing the DAX, and shaping the model underneath it.

Write the DAX

The agent drafts and optimizes measures against your real model, translates SQL or Excel logic into DAX, and explains the pattern so you learn it, not just copy it.

Shape the model

Power BI's project format stores the model as text (TMDL), so an agent can propose and edit tables, relationships, and hierarchies like code you review.

Both skills learn your model, not a textbook one. Because the agent works against your actual schema, the measure it drafts uses your table and column names, and the model edit it proposes fits your relationships. You read the diff the way you would read a colleague's pull request.

Where the human stays

You verify filter context and reconcile the grand total, because a measure that runs clean can still be wrong. And you own the metric definitions and the join semantics. Only you know that "active customer" means something specific to your business, and that a relationship pointing the wrong way quietly doubles a total. The agent writes the code. You certify what it means.

Reports, data prep, and documentation

Three more skills that clear the work you never quite get to.

Build the report page. Describe the page in a brief, or hand over a screenshot of one you like, and the agent builds the visuals and layout as a first draft. You start from something on the canvas instead of an empty page.

Clean the data. Point the agent at a messy table and it profiles the columns, flags broken join keys, and drafts the Power Query (M) steps to fix them. It finds the blank that breaks a relationship before it ever reaches a visual.

Document the model. The agent writes the data dictionary and the measure descriptions straight from the model. This is the work that always gets deprioritized, done in one pass and easy to keep current.

Where the human stays

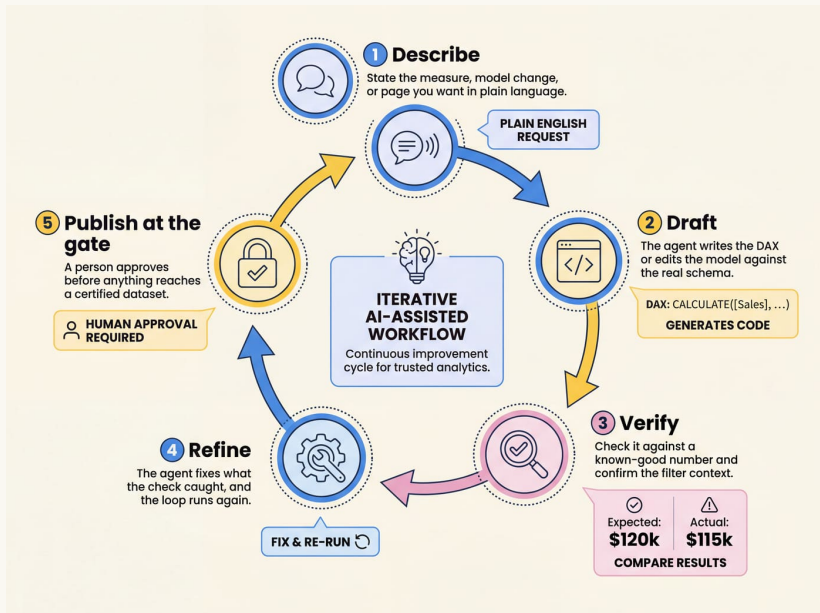
Layout taste and the final read stay with a person. The agent can place twelve visuals correctly and still miss the one chart that tells the story. And the business rules behind the cleaning need you to confirm. Whether a null means zero or means missing is a decision, not a default. Approve the rule, then let the agent apply it.

Audit, performance, and insight

The last two skills close the loop between a model that works and a model you can trust.

Audit and tune. The agent runs your model against best-practice rules, for example Tabular Editor's Best Practice Analyzer, and flags the performance issues that show up in DAX Studio. Then it proposes fixes, from a smarter measure to a column you can drop.

Explain the number. The agent explains why a number moved, for example why margin dropped in one region, and drafts the executive summary to go with it. You get a starting narrative instead of a blank text box, backed by the figures from the model.



Every skill in this part runs the same loop. **Describe, draft, verify, refine, then publish at the gate.** You state what you want, the agent drafts against the real schema, you check it against a known-good number, the agent fixes what the check caught, and a person approves before anything reaches a certified dataset.

PART 3 · HOW AGENTS REACH YOUR MODEL

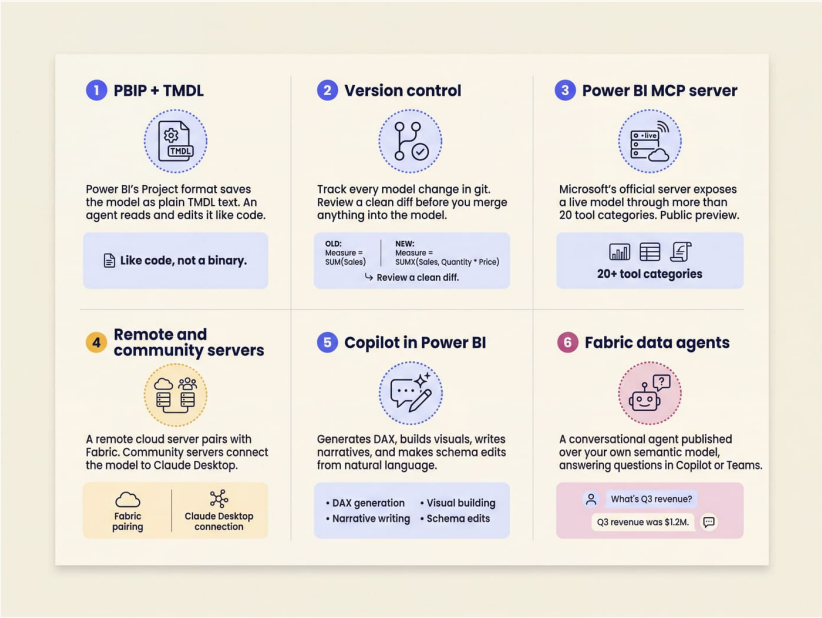
The three doors into your model

An agent does not touch your Power BI model by magic. It comes through one of three doors, and knowing which door you are opening tells you what the agent can see and change.

The text door. Power BI's Project format (PBIP) saves the model as TMDL text files and the report as definition files. An agent reads, diffs, and edits them like code, and you put the whole model under git version control. This is the structural change that lets a general agent work on a real model.

The connector door. A Power BI MCP server exposes the live model to an agent. Microsoft's official Modeling MCP Server has been in public preview since early 2026 and gives an agent tools over tables, measures, relationships, DAX, and more. MCP is an open standard, so the agent connects securely without copy and paste.

The Microsoft door. Copilot and Fabric data agents are built into the platform itself. They live where your data already lives, inside Power BI and Fabric, and need a paid capacity.



The Microsoft-native surface

The Microsoft door is the one most analysts open first, because it lives inside tools you already pay for. It is powerful, and it comes with a licensing reality worth understanding up front.

What Copilot does

Copilot in Power BI generates DAX, creates visuals, summarizes data, writes narratives, and answers natural-language questions about a semantic model. Copilot in web modeling reached preview in 2026. It makes schema updates by plain language, such as renaming tables and columns, creating relationships, and generating measures.

The licensing reality

Minimum: a paid Fabric capacity from F2, or Power BI Premium. A Pro or Premium Per User license alone is not enough to switch Copilot on.	Entry cost: F2 runs roughly two hundred and sixty US dollars a month. Treat all figures as signposts that vary by region.
Free viewers: at F64 and above, people with a free license can view the content. Below that, each viewer still needs a Pro license.	The real prerequisite: a well-curated semantic model. The agent is only as good as the model it reasons over.

Fabric data agents. You build a conversational agent over your own semantic models, then publish it into Copilot Studio or Teams. Fabric IQ, the semantic foundation, is now generally available. Clean measures, defined relationships, and business terms on the right columns are what let the agent reason correctly.

Set up your first skill in an afternoon

You do not need a migration or a budget request to start. Pick the path that matches where you already work, and begin read-only.

Path A: Microsoft-native

Turn on Copilot in the Fabric admin portal if you have the capacity, or build a Fabric data agent over a certified model. Best if you already live in Fabric.

Path B: Claude or general agent

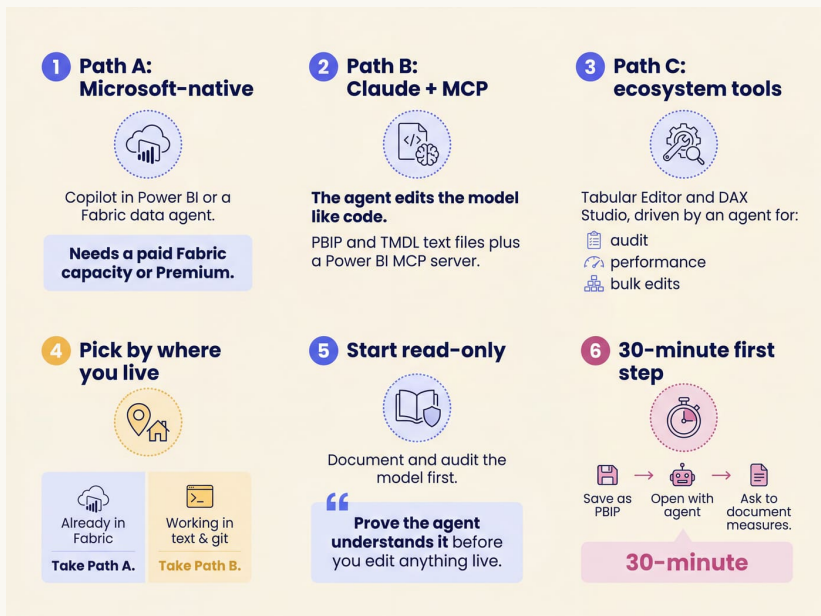
Enable the PBIP format in Power BI Desktop so your model is TMDL text, then point an agent at those files or connect the Power BI Modeling MCP server. Best for git-tracked model work.

Path C: the ecosystem

Drive Tabular Editor and DAX Studio from an agent for audit, performance tuning, and bulk edits across the model.

The 30-minute first step

Low risk, high payoff. Save an existing report as a PBIP project so the model becomes TMDL text. Open it with an agent and ask it to document every measure. You learn exactly how the agent sees your model without touching anything live.



PART 4 · TRUST IT, THEN OWN IT

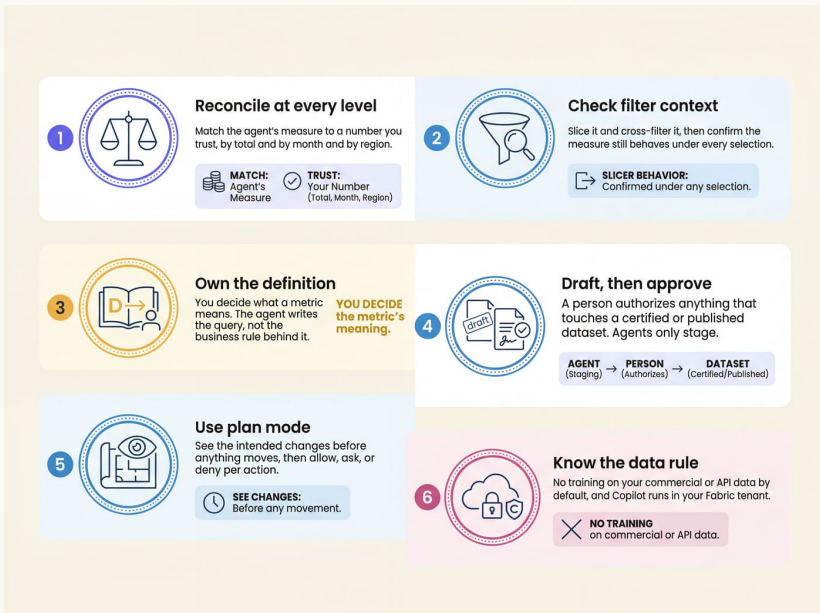
Trust it, then verify

An agent writes DAX faster than you can read it. That speed hides the real risk. The measure that throws an error is harmless, because you see it. The dangerous one returns a plausible number that is subtly wrong, and nobody notices.

A wrong number that looks right is the only dangerous one.

Three habits before you publish

Check it against a known-good number. Take a total you already trust and confirm the agent's measure lands on it, and check it by month and by region too, because a subtle error often reconciles at the total and breaks underneath. **Confirm the filter context.** Slice it, cross-filter it, and watch that it behaves. **Reconcile before anything goes out.**



Then draft, then approve. Agents stage the change. A person authorizes anything that touches a certified or published dataset, and plan mode shows the intended edits first. Name the data rule early too, because it unblocks most security reviews. Anthropic does not train on your commercial or API data by default, and Copilot runs inside your own Fabric tenant.

The analyst's new role

As agents take on the production work, a real part of the job gets automated. The mechanical DAX. The visual dragging. The documentation grind. None of that is where your value lived anyway.

What stays valuable is harder to automate and more worth your time. Data-modeling judgment. Correct metric definitions. Business context. Verification discipline. An agent can write the query, but only you know that "active customer" means something specific to your business.

What gets automated

Mechanical DAX, dragging visuals, and the documentation grind.

What stays yours

Modeling judgment, metric definitions, business context, and verification.

Direct the agent. Do not race it.

The role shifts. You move from writing every measure to directing and verifying the agent that writes them. This touches a very large profession. Power BI is used by tens of millions of people, so the change is not a niche one. The analysts who learn to direct agents will pull ahead of the ones who try to out-type them.

The winning move is clear. Learn to direct and verify agents, not to compete with them on typing speed. That skill is learnable, and it compounds. The next page is where to start.

Where Enterprise DNA fits

We run our own company on exactly this.

Enterprise DNA has spent more than a decade helping people do more with their data, and Power BI is where a lot of that work lives. Everything in this guide is something we practice ourselves. Our own operation runs on a command center of agents, with a person holding the gate, so when a pattern here stops working we feel it before you do.

There are three ways we help, and you can start with any one of them.

Learn it

On-demand courses and live training to get your team fluent with agents, DAX, and modeling an agent can reason over.

Build it

We help you stand up the skills, the connections, and the review habits on your own models.

Run it

We operate the layer alongside your team, the way we run our own command center, and keep it current as the tools change.

Start with the first skill

Pick one report, save it as a project, and ask an agent to document every measure. Thirty minutes, nothing live touched. If you want a faster path, we will find the first three places agents can remove work from your reporting this quarter.

Learn: enterprisedna.co

Email: sam.mckay@enterprisedna.co.nz

Book: a thirty-minute call, no pressure.

The decision: yours.

A NOTE IN CLOSING

You no longer hand-build every report. You describe the outcome, an agent drafts it, and you verify what ships. The analysts who learn to direct agents, and to check their work, will run circles around the ones who kept typing.

FROM INSIDE THIS GUIDE.

Put agent skills to work in Power BI.

ENTERPRISE DNA · AUCKLAND · 2026 · ENTERPRISEDNA.CO