

Microsoft-validated Fabric Workload · also available as a cloud service

Build an AI Readiness Practice on Power BI

BI Pixie scopes the work that takes a client's Copilot and data agents to production, and measures whether AI answered better afterwards. Your consultants deliver it.

The Work Between a Power BI Estate and an AI Answer

Your client's Power BI semantic models already carry the business context AI needs, but it was written for people, not for AI agents.

Making that context unambiguous to an AI agent takes many iterations on a single semantic model, and a client estate holds hundreds of them. Curating an estate at that scale is consulting work that BI Pixie scopes and ranks.

Three Problems Between Your Client and Production

How do you get AI Readiness at scale?

Optimizing one semantic model takes many iterations and tests. A client estate needs that across the mission-critical ones.


A wrong AI answer looks like a right one

An AI agent that lacks context does not report an error. It returns a confident answer that may be wrong, and your client finds out in an executive review.

Stay correct as semantic models change

A new measure or a deleted description can make last month's correct answers wrong, and nothing flags it until a user reports it.

The BI Pixie AI Readiness Loop



What a Partner Gets Out of It

- ◆ **The remediation revenue is yours in full.** BI Pixie does not deliver consulting and does not compete with you on services. BI Pixie names the work, and your consultants do it.
- ◆ **A proposal you can price.** BI Pixie returns a ranked list for every semantic model, naming what is holding AI back and which fix is worth making first.
- ◆ **A free first look, then a plan sized to the work.** A first assessment of up to 500 semantic models is read-only and free on every plan. A project-scale before and after benchmark sits on the Pro plan.

Where BI Pixie Fits in Your Projects

- ◆ **A Fabric or Power BI migration.** Beside retire, repoint, refactor and rebuild, your consultants take a fifth decision: which semantic models to prepare so Copilot and Fabric data agents answer correctly. A report nobody opens is an easy retirement. Its semantic model may be worth keeping.
- ◆ **A Copilot rollout.** Your client learns whether their semantic models can serve Copilot accurately before their executives do.
- ◆ **A managed service after the project closes.** On the Enterprise plan BI Pixie assesses semantic models on a schedule your client sets.

Three Ways Partners Work With Us

Resell BI Pixie

Your firm carries the license and bills your client for it.

Sell alongside us

Your client buys BI Pixie directly, and your firm delivers the work.

Refer

Your firm introduces the client and stays out of the transaction.

We agree which one applies to your firm in the partner agreement.

Bring Us Into Your Next Client Conversation

- Free assessment of up to 500 semantic models
- Read-only, and nothing is written without approval
- Talk to us at bipixie.com/partners



app.bipixie.com