

Stop Reading Dashboards, Start Asking Questions

Process Mining + Copilot Studio: ask your process anything

Elliot Margot - Microsoft MVP (M365 Copilot & Copilot Studio)

17 June 2026

At a glance

Platform	Power Automate Process Mining
Connector	Process Mining MCP (Preview) - 9 tools, generative orchestration ON
Agent	Copilot Studio - Process Insights
Processes	Procure-to-Pay and Power Automate Flow Health (both fully synthetic)
Runtime	About 8 to 10 minutes - 4 questions, live

Before you start

- ☐ Agent published and tested in Copilot Studio, on the environment that holds your processes.
- ☐ Process Mining MCP (Preview) connected; generative orchestration ON.
- ☐ get_processes returns BOTH processes: Procure-to-Pay and Power Automate Flow Health.
- ☐ Test pane open, or a published channel ready (Teams or Microsoft 365 Copilot).
- ☐ Internet is live - every MCP call hits the cloud engine.
- ☐ Synthetic data only on screen. No client data, ever.

Open with this

Dashboards are passive: you read the charts and go hunting for the why. This flips it. You ask in plain language and get grounded answers straight from the live process data - every number traceable, never a guess.

The demo - four questions

1. "Where is the biggest bottleneck in Procure to Pay?"

Agent calls get_bottleneck_analysis, then get_edges_with_metrics

You should see Approve Purchase Order is the bottleneck. The transition Revise PO -> Approve PO waits about 6 days 5 hours across 447 cases (14.9% rework); Create PO -> Approve PO about 5 days 20 hours.

Say Watch it pivot on its own: activity durations look flat, so it moves to the edge wait-times to find the real delay.

2. "What is driving the long cycle times?"

Agent calls `get_correlation`

You should see One vendor, Globex, dominates: Duration Influence +4.22, about 5x any other vendor. The order Amount barely matters.

Say *This is the WHY, not just the WHERE: a correlation the engine computes, not a chart you have to interpret.*

3. "Are there rework loops?"

Agent calls `get_variants_with_metrics`

You should see About 15% of cases follow reject -> revise -> re-approve before they complete.

Say *Rework is hidden cost. The agent quantifies it in one sentence.*

4. "Which Power Automate flow is failing the most, and where does it break?"

Agent calls switches to the Power Automate Flow Health process (`get_processes`, then the bottleneck and edge tools)

You should see Invoice Sync, 373 failures (about 3x the next worst), breaking at POST invoice (HTTP) with Throttled (429) - even after a retry.

Say *The meta moment: we are mining our OWN failing flows, not a textbook dataset.*

Land this point

The point. Every number you just saw came from the Process Mining engine through MCP, not from the model. The agent reads the data; it does not imagine it. That is the anti-hallucination guarantee.

If something goes wrong

- **`get_processes` returns nothing** - check the MCP connection and that you have Process Mining Contributor or Viewer access to the environment.
- **Calls take 10 to 30 seconds** - narrate while it thinks: it is querying the live engine right now.
- **An answer looks flat or generic** - ask it to look at the transitions or edges, or name the process explicitly.
- **Keep the environment clean** - `get_processes` lists every process in the env; remove anything you do not want on screen.

From dashboards to dialogue.

Grounded, conversational analytics anyone can ask.

Resources

- **Talk page and slides** e-margot.ch/community/talks/stop-reading-dashboards
- **Demo repo (data, generators, agent prompt)** github.com/OwnOptic/stop-reading-dashboards-demo
- **Build it yourself (Microsoft Learn)** learn.microsoft.com/.../process-mining-mcp-create-cps-agent