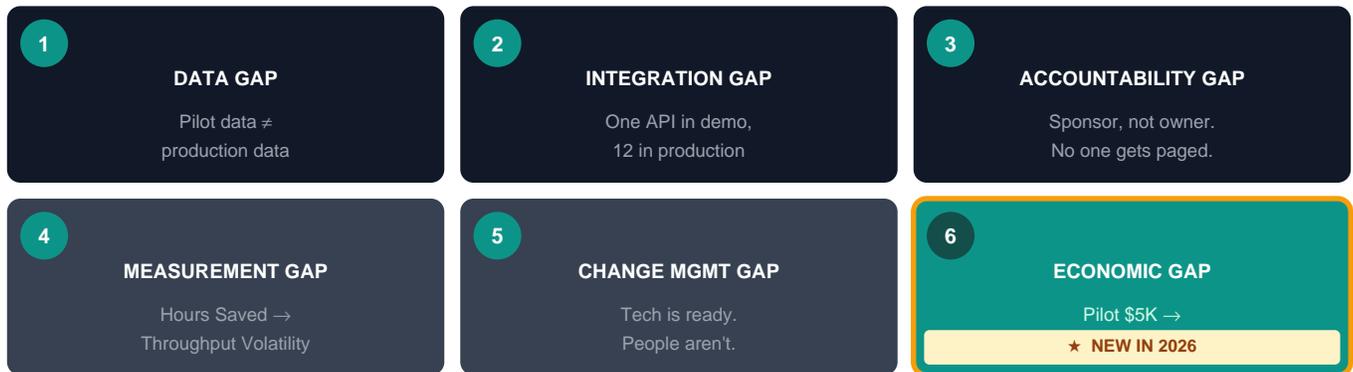


## SIGNATURE FRAMEWORK NO. 3

# The Pilot Trap™

The six gaps that kill enterprise AI between demo and deployment

***"The enterprise AI graveyard is full of impressive pilots. They worked. The demos were compelling. The sponsors were excited. And then... nothing. The Pilot Trap maps exactly why — and exactly which gap killed yours."***



Most failed agentic initiatives don't fail at Layer 1 (capability). They fail at the crossing — when a working pilot meets the real complexity of production: messy data, 12 integrations, no named owner, a measurement framework that can't capture scale, and a cost model that was never stress-tested at 100x volume.

## THE SIX GAPS — QUICK DIAGNOSTIC

GAP	THE TRAP	THE DIAGNOSTIC QUESTION
1 Data	Pilot uses clean, curated data. Production data is messy, distributed, inconsistent.	Is the pilot data the same source, format, and completeness as what the deployed agent will actually see?
2 Integration	One API in the demo. 12 integrations in production, many with 60-day security review cycles.	Have all required integrations cleared your security approval process? What's the longest approval cycle?
3 Accountability	Enthusiastic sponsor. No named owner post-deployment. Agent makes a mistake — who gets called?	Who is the named human accountable for this agent's decisions after deployment? What happens when it makes a mistake?
4 Measurement	Measuring Hours Saved. Should be measuring Throughput Volatility: performance during 500% volume spikes.	Can it handle a 500% volume spike at the same cost-per-decision and the same accuracy?
5 Change Mgmt	Technology ready. People still manually double-checking every decision. Adoption at 12%.	For every affected role: what's their new responsibility, and how will they be evaluated?
6 Economics ★	Pilot: \$5K for 10 users. Production: potentially \$500K+ for 10K users. Unit economics were never modeled.	What is your cost-per-decision at 100x volume? Has your CFO reviewed that number?

## THE INFERENCE RECKONING (GAP 6)

Gap 6 is the newest and fastest-growing source of deployment failure in 2026. The fix is **unit economics from day one**: design the pilot to produce a cost-per-decision figure at P50 and P95 volume. Architecture decisions that protect unit economics: cascade model selection (small models for routing, large for complex decisions), aggressive caching on repeated patterns, batch processing where real-time isn't required, and context window management (retrieve only relevant passages, not entire documents).

## HOW TO USE THIS FRAMEWORK

- **Keynote:** "How many of you have a pilot that worked but never deployed?" Most hands go up. "I'm going to tell you exactly which gap killed it."
- **Customer conversation:** Use as a diagnostic, not a pitch. Identify the buyer's specific gap before they hit it. That moves you from vendor to advisor.
- **Thought leadership:** Counter-narrative to "AI is transforming everything" — the problem isn't that pilots fail. It's that they *succeed*, and we haven't built the crossing infrastructure.