



NETSURIT

The Board Has Stopped Counting Pilots

**And started counting dollars: how AI
multiplies productivity and where it does not**

by Louis de Klerk - Chief Information Officer

June 2026

The Question On Every Board Agenda



Two years ago, "have you tried ChatGPT?" was an interesting question at a board meeting. Eighteen months ago, "what is our AI strategy?" became a mandatory one. Today, the question has sharpened to its hard form:

"How is AI going to multiply our profitability while lowering our costs, and when?"

The answers companies are giving are not landing well. The data, as of mid-2026, is sobering:

- MIT puts the enterprise GenAI failure rate at **95%**. These are projects that did not produce measurable financial returns. (MIT NANDA, "The GenAI Divide: State of AI in Business 2025" — 150 leader interviews, a 350-employee survey, and analysis of 300 public deployments.)
- **97%** of enterprises still struggle to demonstrate business value from their early generative AI efforts. Only **29%** see significant ROI from generative AI, and **23%** from AI agents. (WRITER 2026 "AI Adoption in the Enterprise" survey of 2,400 knowledge workers, including 1,200 C-suite executives.)
- Only **3.3%** of Microsoft's 450 million Microsoft 365 commercial subscribers have signed up for the paid Copilot add-on, despite Microsoft's \$37.5 billion quarterly AI investment. (Microsoft's Q2 FY26 earnings disclosure.)
- **74%** of companies using AI tools cannot show tangible business value despite widespread adoption claims. (Industry analysis aggregating BCG and Deloitte findings.)
- **73%** of CEOs report stress or anxiety about their company's AI strategy, 39% describe it as "crippling," **61%** fear losing their job if they fail to lead the transition, and **75%** admit their company's AI strategy is "more for show" than actual internal guidance. (Same WRITER 2026 survey.)

The Question On Every Board Agenda



Microsoft's own Jared Spataro, who runs Modern Work and Business Applications and is therefore the most senior person in the world responsible for selling Copilot, said the part out loud:

"Even when Copilot improves efficiency by 20–30% in tests, connecting that to tangible ROI is tricky. In knowledge work, being '30% more productive' doesn't automatically show up on the bottom line. If an engineer writes a report faster, it doesn't directly create 30% more profit."

That sentence is the entire argument of this article. The AI tools your team uses today produce real individual wins, with measurable hours saved, and they are not converting to enterprise profitability. The board's frustration is not a perception problem; it is grounded in real data.

This article is about three different shapes of AI investment: what each actually delivers, how much each actually costs, and how to spot vendors who'll claim the most valuable shape without delivering it.

Three Shapes Of AI Investment




Only one of them answers the board's question



The three shapes:

- 1 Personal AI** makes one person faster. Microsoft Copilot in your Word and Outlook, ChatGPT and Claude.ai for chat, Cursor and GitHub Copilot for code. Well established. Modestly priced. Modest in what it changes about the business.
- 2 Team AI** coordinates a team's chat-anchored work. Claude Cowork (Anthropic, April 2026), ChatGPT Workspace Agents (OpenAI, replaced Custom GPTs April (2026), Claude.ai Projects with Live Artifacts. Growing in popularity. More nuanced pricing than the seat fees suggest.
- 3 Process AI** encodes and runs a business process. The new category. The one that is most likely to satisfy the board's demands. The one every AI vendor on earth is about to claim they have.

A buyer's checklist follows.

 Personal AI	Faster individuals
 Team AI	Better coordination
 Process AI	Business operating leverage

Why Personal AI Hits A Ceiling On Bottom-Line Impact



Personal AI is Microsoft Copilot in your Word and Outlook, ChatGPT and Claude.ai for chat, Cursor and GitHub Copilot for code. Pricing is well-established and modest: Microsoft 365 Copilot Business is \$18/user/month (promotional through 30 June 2026, \$21 after); the Enterprise add-on is \$30/user/month. ChatGPT Plus is \$20/month; Claude Pro is \$20/month.

Forrester models project SMB Copilot ROI at 132% to 353%, with savings of about nine hours per user per month, roughly 2.25 hours per week (Forrester, "Total Economic Impact of Microsoft 365 Copilot for SMB", October 2024). Those numbers are real in the sense that the hours genuinely are saved. **The question is what happens to those hours.**

In knowledge work, saved hours have four destinations, and only one of them shows up on the P&L:

- 1 Same headcount, slightly more output.** A salesperson drafts proposals 30% faster. They send four proposals a week instead of three. If there's unmet demand, this captures incremental revenue. For most SMBs, sales velocity is gated by lead generation, decision cycles, or capacity to deliver — not drafting speed. The extra proposals mostly do not translate into more sales.
- 2 Slack time absorbed by everything else.** The accountant who saves an hour on email spends it on the slightly-overdue training module, the longer coffee, the meeting that runs five minutes late. Parkinson's law applies as forcefully to knowledge work as to government departments.
- 3 Higher quality on the same output.** The report that took three hours now takes two, and the engineer spends the third hour polishing it. Genuinely better work — but it doesn't show up at the bottom line, because the firm wasn't being penalized for the previous lower quality level.

Why Personal AI Hits A Ceiling On Bottom-Line Impact



- 4 **More meetings, more email.** Faster drafting of internal communications begets more internal communications. Some studies suggest GenAI is increasing email volume — something I see personally.

None of those destinations multiply profit. They diffuse value, raise the quality floor of work, and slightly reduce friction. All real. None of them visible to the CFO.

Where Personal AI does show bottom-line impact is narrow and specific: software-engineering throughput when the engineering team is the bottleneck on revenue; customer-support deflection when a chat agent reduces ticket volume that was being staffed; high-volume creative production where output is metered. Outside those bands, Personal AI is doing what it was designed to do, which is making individuals faster — and what it was designed to do is **not the thing the board is asking for.**

The board is asking for operating leverage: more output without more headcount, or the same output with less headcount. Personal AI does not produce operating leverage. It produces personal-productivity leverage, which is a different and more "fluffy" thing.

-  **This is the productivity paradox in one sentence: individual wins are real and measurable, and they're not translating to business value.**

Until that distinction is truly understood at board level, every Copilot rollout will read like the previous one.

What Team AI Buys You



Team AI is the assistant-shaped agentic tier: Claude Cowork (Anthropic, April 2026), ChatGPT Workspace Agents (OpenAI, replaced Custom GPTs April 2026), Claude.ai Projects with Live Artifacts. Pricing is more nuanced than the seat fees suggest:

- **ChatGPT Business** is \$25/user/month (annual); Enterprise is approximately \$60/user/month. Workspace Agents are free during research preview through 6 May 2026, then move to credit-based pricing — meaning usage on top of the seat fee.
- **Claude Team Standard** is \$20-25/seat/month; Team Premium is \$100-125/seat/month (includes Claude Code); Enterprise is custom-priced. Crucially, Anthropic Enterprise seat fees explicitly do not include token usage. Every chat, every Cowork action, every Claude Code run is billed at API rates on top.

Total cost-of-ownership is seat fee + usage, not just the seat fee. For a heavy team, the usage line can equal or exceed the seat line.

What Team AI raises the floor on:

- **Team coordination.** Shared Projects with multi-user editing and version history mean a partner, a manager, and a staff accountant all work against the same document set in the same agentic context.
- **Process consistency across staff.** A Skill or plugin packages an instruction set once, with optional Python and templates. Everyone using that Skill gets the same starting point, and Team AI makes sharing these Skills easy.
- **Scheduled recurring tasks.** ChatGPT Workspace Agents schedule daily, weekly, or monthly. Claude Cowork has Dispatch. Monday-morning briefings, weekly compliance digests, monthly close prep: these become unattended.

What Team AI Buys You



- **Connector access.** Drive, Gmail, SharePoint, Salesforce, FactSet, DocuSign. The agent reaches your business apps without manual copy-paste.
- **Lower-friction onboarding.** A new hire opens the workspace, sees the firm's published Skills and Projects, and produces immediately.

Where Team AI still stalls:

- **Process unit economics don't change much.** A Workspace Agent that drafts engagement letters more consistently is still operating at human-scale throughput, with a human in the chat session for each engagement. Headcount-per-engagement doesn't move materially.
- **Audit posture is the chat transcript.** When a regulator asks "why did this engagement letter use that liability cap?", the answer points to a chat session, often concluding in "the LLM chose it, but we don't know why".
- **The agent doesn't run while the team sleeps.** Most Team AI products are still chat-anchored. Even when scheduled, they execute and surface results into a chat surface someone reads in the morning. That's better than not running, but it isn't a process executing overnight against your data.
- **The work product belongs to the chat, not the firm.** Each user's chat history accumulates context the next user can't see.

Realistic Team AI impact for a typical **SMB: 10–25% productivity uplift on team-process work, plus measurable consistency improvements, plus a reduction in onboarding time.** That's worth real money. But it is not "multiply profitability." It is "raise the floor."

For the board's question, Team AI is a better answer than Personal AI. **It is not yet the answer.**

What Process AI Unlocks

Encoding processes, not amplifying people



“The work where you'd rather encode a process than amplify a person.”

That is the line. It is the only line that matters for understanding Process AI.

Personal AI amplifies a person. Same person, faster. Team AI amplifies a team. Same team, more consistent, with shared context. Process AI **encodes the process itself**. The process runs against the data, with audit-trailed checkpoints where humans decide, producing structured deliverables, on a schedule, at cloud throughput. The person's role shifts from doer to reviewer-and-decider.

The mathematical difference is substantial. If Personal AI makes a bookkeeper 30% faster at coding transactions, the firm needs slightly fewer bookkeeper-hours per client. If Process AI encodes the bookkeeping methodology, the transactions are automatically processed overnight, and the bookkeeper handles five times as many clients as the firm's bottleneck allows. **That ratio is what shows up at the bottom line.**

Three patterns recur:

The monthly close that runs overnight. A small accounting firm's month-end close currently consumes the better part of a week per client: bank reconciliations, journal review, intercompany matching, accruals, trial balance review. Encode that as a Process AI workflow. The skill mounts the client's bank-statement and ledger shelves, runs reconciliations (matching transactions deterministically, no LLM in the hot path), drafts journal entries, surfaces exceptions to a review queue, runs overnight on the first of every month, and presents the controller with a fully-staged close at 8 AM.

The controller reviews exceptions, approves at the gate, and the books close by lunch. **The headcount-per-close ratio drops by ~70%; the client capacity per controller rises by ~3x.**

What Process AI Unlocks

Encoding processes, not amplifying people



The invoice queue that triages itself. A 200-employee distributor processes 3,000 vendor invoices per month. The AP analyst spends most of their day on transcription, duplicate-checking, and policy verification. With Process AI you encode the eight-specialist invoice review pattern as a workflow. Triage runs continuously as files land in the shared folder. The analyst opens a review queue every morning where 90% of items are pre-approved straight-through and the remaining 10%, the genuine exceptions, arrive on their desk already sorted by exception type, with the institutional vendor memory attached, the verdicts of all eight AI specialists visible at a glance, and the proposed disposition pre-filled. The analyst's day shifts from transcription to judgment. **The same headcount processes 4–5× the invoice volume.**

The discovery sentinel that wakes weekly. A specialty manufacturer has 5,000 active customer accounts and tracks regulatory compliance changes across the eight states they ship to. Currently a compliance manager spends Monday morning reading regulatory bulletins. Encode an AI sentinel: scheduled wakeup Sunday night, scrapes regulatory feeds, cross-references against the customer base, surfaces the seven customers affected by the new South Carolina rule and the two affected by the Massachusetts amendment, drafts the customer notification in the firm's voice, presents a Monday-morning briefing with one click to send. The compliance manager's role shifts from researcher to reviewer; the firm catches changes it would have missed; the audit trail is structured rather than anecdotal.

In each pattern, the human decision points are preserved (gates 1, 2, 3) but the unsupervised processing volume between gates is what changes. **That is operating leverage. That is what the board is asking for.**

The honest catch: encoding a process is real work. It is not a Copilot rollout. It is a discrete engineering investment per process. Which brings us to the cost shape.

The Investment Shape

Three different cost models



The three categories have fundamentally different cost models, and the board needs to understand which one it is approving.

Personal AI is opex per seat. \$18–\$30 per user per month for Microsoft 365 Copilot, \$20 for ChatGPT Plus or Claude Pro. Low upfront cost. Distributable across the whole staff. Low effort to deploy. Low and diffuse impact. The right mental model is "raise the floor of personal productivity." Genuinely worth doing, and not in tension with anything else.

Team AI is opex per seat plus usage plus integration effort. \$25/seat for ChatGPT Business or Claude Team Standard at the SMB-realistic tier; \$60/seat for Enterprise. Plus usage credits or token charges layered on top. Plus a discrete project to wire up Skills, connectors, and shared Projects. For a 50-person firm deploying selectively to managers and above, total year-one cost lands around \$30–70K depending on usage intensity. Medium and team-level impact: better coordination, better consistency, faster onboarding.

Process AI is platform investment plus per-process engineering. The platform itself runs in the firm's own Azure subscription. Infrastructure costs (compute, storage, model API spend) typically in the low five figures per month for an SMB at modest scale (illustrative; varies materially by volume). Per encoded process, think of it as a small product-engineering project: typically 4–10 weeks of methodology owner plus one or two engineers to encode a real line-of-business process the firm currently runs by hand. Transformational unit-economic impact on the encoded process.

The right mental model for Process AI is **capital expenditure on a process, not opex on an AI seat.** You are buying down the recurring cost of running that process by paying once to encode it. The payback math is:

$$(\text{annualized hours saved per run} \times \text{runs per year} \times \text{loaded hourly cost}) - (\text{encoding cost})$$

The Investment Shape

Three different cost models



For a process that runs daily, weekly, or per-client (a monthly close runs ~12 times per year per client × number of clients), the math is usually obvious within 12 months. For a process that runs three times a year, it usually isn't worth encoding. That's a feature, not a bug, of the cost model. **Process AI is for the work you do often.**

Team AI and Process AI are not in competition. They serve different pieces of the work portfolio. A well-run SMB will use both: Team AI for the long tail of team-coordination work, Process AI for the three-to-eight processes that run often enough to be worth encoding.

Shape	Cost model	Year-one cost (50-person firm)	Impact
Personal AI	Opex per seat. \$18–\$30/user/month.	~\$11K/year for 50 seats at \$18.	Raises the floor of personal productivity. Essentially zero P&L impact.
Team AI	Opex per seat + usage credits + integration project.	~\$30–70K depending on usage intensity.	10–25% productivity uplift on team-process work. Net positive, not transformational.
Process AI	Platform investment + per-process engineering (4–10wks each).	~\$180K engineering + ~\$60K/year infra for three processes.	Headcount-per-process drops 60–80%; client capacity rises 2–5x.

A Worked Example

A 50-person US accounting firm



Concrete numbers, with caveats. The dollar figures below are illustrative, derived from typical reference deployments and US labor market norms — not from any specific firm's books. A regional CPA firm: 8 partners, 25 staff and senior accountants, 17 admin and support. Mix of monthly bookkeeping clients, quarterly tax-prep clients, and annual audit clients. Annual revenue ~\$11M. Loaded cost per accountant ~\$130K; per admin ~\$70K. Personnel cost ~\$5.5M.

Personal AI. Microsoft 365 Copilot Business across the firm at \$18/seat/month → ~\$11K/year for 50 seats. Each accountant saves ~2 hours per week on email, document drafting, and meeting prep. P&L impact: essentially zero. Saved hours don't convert to billable hours or client capacity. Worth doing for floor-raising and quality-of-life; not the answer to the board's question.

Team AI. ChatGPT Business at \$25/seat for 15 manager-and-above seats → ~\$4.5K/year, plus ~\$15–25K of usage credits (variable), plus a \$40K integration project for Skills covering engagement-letter drafting, IRS-notice response, and audit-workpaper review. Year-one cost: ~\$60–70K. Realized impact: engagement letters draft 60% faster and more consistently; new-hire onboarding drops from 8 weeks to 5; the firm catches deadlines it would otherwise have missed. Estimated annualized benefit: ~\$120–180K. **Net positive, real, but not transformational.**

Process AI deployment. Platform installed in the firm's Azure tenant. Three Experiences encoded over six months, in priority order:

- 1 Monthly bookkeeping close.** Runs overnight against client bank-statement and ledger shelves; staff role becomes review-and-exception. Encoding cost: ~\$80K. Recurring marginal cost: model API + compute, ~\$15/client/month.
- 2 AP invoice review** for clients with high-volume AP. The eight-specialist pattern adapted for client-level processing. Encoding cost: ~\$60K. Reuses the close-process Library substrate.

A Worked Example

A 50-person US accounting firm




- 3 **Year-end 1099 generation and reconciliation.** Annual but high-effort. Encoding cost: ~\$40K.

Total year-one engineering: ~\$180K. Platform infrastructure: ~\$60K/year.

Realized impact: the firm's bookkeeping practice can absorb a **2–2.5× increase in client count without adding bookkeeper headcount**; existing bookkeepers move into review-and-advisor roles at higher billable rates; the firm captures market share its competitors can't service. The 1099 season that previously consumed three weeks of staff time and partner sign-offs runs in three days of review work. AP-heavy clients see month-end turnaround drop from two weeks to two days.

If the firm successfully takes on 30 net new bookkeeping clients at ~\$2,400/month each on the freed capacity, that's **\$864K of incremental annual revenue at near-100% gross margin** (the workflow runs at marginal compute cost). Net of the encoding investment in year one, the bookkeeping process alone pays for the entire Process AI program. By year two, the program is a profit center.


-  These figures are illustrative and the practical benefit will vary. The shape is what matters: **Process AI economics are nonlinear in the throughput axis.** It is the only one of the three categories that gives a board a credible answer to "multiply profitability while lowering costs."

Other SMB Shapes

Same pattern, different processes



The pattern repeats; the encoded processes differ.

- **A 75-person specialty manufacturer.** PO processing, supplier onboarding, quality-incident analysis, and sales-tax nexus tracking. Estimated freed capacity: 25–40%, mostly captured as avoided headcount against a planned hire.
 - **A 30-attorney boutique law firm.** Intake conflict checks, first-pass contract review, discovery production, privilege flagging, and redaction proposals. Routine-work throughput per associate: ~1.5–2x; partner time freed for complex matters.
 - **A 20-provider primary-care practice.** Prior-authorization letters, claims-denial review, reason-coded resubmission drafts, and patient document chase for visit prep. Net effect: 3–4 fewer FTE in billing and administration.
 - **A 60-person professional-services firm.** Proposal generation, project-status sentinels, and onboarding-kit assembly. Capacity to bid ~40% more proposals with the same business-development headcount; status reporting becomes free.
-  **Common thread:** identify the three-to-eight processes the firm runs most often, encode those, leave the rest to Personal and Team AI.

Team AI transitioning into Process AI. As these platforms evolve, the boundaries will not stay clean. Workspace-agent tools are already moving toward the lower end of Process AI: shared agents that use connectors, follow team instructions, run on schedules, request approvals, and push work across business tools. Some will absorb legitimate process use cases, especially where workflows are light, data volume is modest, and audit requirements are low.

The distinction is not whether the vendor says “workflow,” “agent,” or “automation.” **It is architectural.**

How To Spot Real Process AI

When every vendor claims to have it



When a category becomes valuable, every adjacent vendor claims they already have it. "Agent" used to mean something specific; today every chatbot vendor calls their product an agent. The same thing is about to happen to Process AI. Cowork's marketing already speaks of "customizable plugins that encode institutional knowledge and workflows" — that's two-thirds of the Process AI pitch, with very different mechanics underneath. Every SaaS vendor whose product touches a workflow will, by Q4 2026, have rebadged their dashboard as "Process AI" (or whatever term the market settles on).

Marketing copy will not survive the seven questions below. They are demonstrable, falsifiable, and architectural. A vendor either can show you each one in a real demo, or they can't.

1 Where does work execute when the office is closed?

Real Process AI: the workflow wakes itself at 02:00, runs against a file share, writes results, hibernates. Show me the scheduled run from last night and the structured output it produced before anyone arrived.

Imitation: "scheduled task" that pings a human to come act on it. The work waits for the human in the chat — which the human should then download or copy out; the human doesn't wait for the work.

2 Where does the work product physically land?

Real: in your file share, your database, your ERP / CRM, your tenant. Accessible from Azure Storage Explorer, from SharePoint, from a System of Record.

Imitation: in the vendor's UI. To get it out, you export. To share it, you grant access to the chat.

How To Spot Real Process AI

When every vendor claims to have it



3 Show me what happens if we crash at minute 25 of a 30-minute process.

Real: the workflow resumes from the last completed phase boundary because every prior stage's output is on disk. Show me the resume from the actual log.

Imitation: it starts over, or produces a half-finished output nobody can use. "We retry" is not the same as "we resume."

4 Show me the deterministic code that runs.

Real Process AI ships Python and SQL inside the workflow. The LLM isn't in the hot path of the heavy computation. The AP arithmetic specialist is `sum(line_items) == total`, not "ask the LLM." Show me the deterministic code. Read it to me.

Imitation: "the LLM handles all of it." That means inference cost scales with volume, audit defensibility is whatever the model felt like saying, and the workflow falls over the moment the input exceeds the context window.

5 Show me where the agent's memory lives, and read it to me.

Real: markdown files I can read, version-controlled, captured during nightly curation. Open the memory file and show me what was added last week and which session produced it.

Imitation: opaque vector embeddings; or "session memory" that resets; or a knowledge graph inspectable only through a chat query. If you can't read it, you can't audit it. Even worse: "we build up memory for you, but won't let you see or control it."

How To Spot Real Process AI

When every vendor claims to have it



6 Show me the audit trail of one decision made six months ago.

Real: the structured artifact captured at the moment of decision: reason code, specialist verdicts, the on-file institutional state at that moment, the human action and timestamp. One click from "approve this invoice" to the full evidence chain.

Imitation: a chat transcript reconstructed from logs. Maybe.

7 Where does my data physically live, and who else's data sits next to it?

Real: in my Azure subscription, my tenant, my encryption keys. The vendor's code runs against my data; my data does not travel to the vendor.

Imitation: in the vendor's multi-tenant SaaS, with whatever data-residency and isolation assurances they can offer in the contract.

If a vendor cannot demonstrate all seven, what they have is Process AI in name only: a Personal or Team AI product wearing the new category label. **That product may still be valuable for what it is. But it is not what your board is asking for.**

How To Position This To A Board Today



If you're the person standing in front of a board being asked "what is our AI ROI?", here's the framing that works:

- 1 Set the right expectation about Personal AI.** It raises the floor; it doesn't multiply profitability. Budget it as opex like Office 365. Don't promise the board operating leverage from a Copilot rollout. The board has seen enough data to know better.
- 2 Deploy Team AI pragmatically.** Pick two or three high-leverage skills (engagement letters, IRS notices, monthly status reports) and encode them at the Team AI tier. Expect medium impact. Don't expect it to change the unit economics of process work.
- 3 Pick three processes to encode in Process AI.** This is the part the board actually wants to hear. Pick them by frequency-of-execution $\hat{\Lambda}$ current-cost-per-execution. The monthly close. The AP queue. The compliance sentinel. The proposal pipeline. Whichever three the firm runs most often. Budget the platform installation plus per-process encoding. **Expect a 12-to-18-month payback per process.** After that, the unit economics change permanently.
- 4 Evaluate Process AI vendors with the seven questions, not their slide deck.** As Process AI becomes recognized, vendor claims will explode. Most will fail at least three of the seven. Hold the line. If you don't find the product you need, wait a few months — this space is rapidly evolving and will likely only reach maturity in 2027.
- 5 Be honest about what's required.** Process AI is not "buy a license and watch productivity jump." It is "encode the firm's methodology into software." You need a methodology owner paired with engineering capacity. Firms that won't fund that will spend the next three years explaining to their boards why their AI strategy is "more for show."

Closing

Why this matters for SMB and mid-market firms



The good news for SMBs: Process AI economics scale better at SMB size than enterprise size. Encoding three processes is tractable for a regional CPA firm; it is not for a Fortune 500 finance function. SMBs that move first in their vertical create durable cost advantages slower competitor cannot easily match.

The work where you'd rather encode a process than amplify a person.

That sentence is the investment thesis for Process AI. Personal AI fits the long tail of unique work, where the bottleneck is the speed of the human's pen. Team AI fits coordination, shared context, and consistency. Process AI fits repeated work, where the methodology is the asset and the constraint is throughput per person. **There, encoding the process is the only thing that produces operating leverage.**

Boards are right to be skeptical of Personal AI productivity claims. They are right to be cautiously optimistic about Team AI. And they are right to ask whether something can finally deliver the operating leverage they were promised. **Process AI is that something.**

Over the next 18 months, vendors will claim it whether or not their architecture supports it. The firms that internalize the distinction now, and evaluate vendors on the seven questions rather than the marketing, will be the ones whose 2027 P&Ls answer the board's question.

If your firm runs three or more processes by hand that would be better encoded than amplified — month-end close, AP review, compliance monitoring, proposal generation, prior authorization, intake screening, or similar repeatable workflows — the next conversation is which ones, in which order, with what payback.

That is the conversation that turns "more for show" into operating leverage your 2027 board meeting can point to.

Sources

Verified against cited sources in May 2026



Failure-rate and ROI data

- MIT NANDA initiative, The GenAI Divide: State of AI in Business 2025. Primary research behind the 95% figure. Coverage: Fortune, Healthcare IT News, AI Magazine, Digital Commerce 360.
- WRITER, Enterprise AI adoption in 2026: Why 79% face challenges despite high investment. Primary source for the 97% / 29% / 23% / 73% / 39% / 61% / 75% figures from a 2,400-respondent survey fielded Dec 2025–Jan 2026.
- The Register, Microsoft reveals just 3.3% of Copilot Chat users pay for it. Based on Microsoft's own Q2 FY26 earnings call.
- Petri, Why Microsoft Copilot Adoption Is Lagging: The ROI Dilemma. Source for the 74% figure and the Spataro quote.
- Forrester, The Total Economic Impact of Microsoft 365 Copilot for SMB (October 2024). Source for the 132%–353% ROI projection and 9-hours-per-user-per-month saving.
- Forrester / Microsoft, The Total Economic Impact of Microsoft 365 Copilot (general). Companion enterprise-tier study.

Pricing references (verified May 2026)

- Microsoft 365 Copilot Plans and Pricing. Microsoft 365 Copilot Business at \$18/user/month (promotional, \$21 after 30 June 2026); Enterprise add-on at \$30/user/month.
- Plans & Pricing | Claude by Anthropic. Pro \$20/month; Team Standard \$20-25/seat; Team Premium \$100-125/seat; Enterprise custom.
- ChatGPT Pricing | OpenAI. Business at \$25/user/month (annual); Enterprise approximately \$60/user/month.
- Introducing workspace agents in ChatGPT | OpenAI. Confirms the move to credit-based pricing on 6 May 2026.
- Anthropic Enterprise plan FAQ. Confirms that Enterprise seat fees do not include token usage.

Note on numbers: survey statistics shift quarter to quarter; the figures above were verified against the cited sources during May 2026. The worked-example dollar figures in the 50-person CPA firm section are explicitly illustrative and derived from typical reference deployments and US labor-market norms, not from any specific firm's books.