AN EXECUTIVE'S GUIDE TO BUDGETING FOR AI

The AI Budget Trap

Why companies are blowing through their AI budgets, and what the few who get a real return are doing differently.

"Budgeting for AI is not a software implementation. It is a workforce expansion. Missing that is the single most expensive mistake an executive can make."

Dave Albertson CEO, ManageAI

manageai.io

Where we are today

A year ago, “AI budget” wasn’t a line item. There was no benchmark, no playbook, no peer to call and ask how they pulled it off. Most companies spent a little on a few licenses, filed it under software, and moved on. Then the headlines started.

In December 2025, Uber rolled out Claude Code to its roughly 5,000 engineers. By April, four months later, the company had burned through its entire 2026 AI budget. Adoption had jumped from about a third of engineers to more than 80%, some of them running up $2,000 a month in usage. Leadership was, in their own words, “back to the drawing board,” capping spend at $1,500 per person and openly admitting they couldn’t draw a clear line between the bill and a better product for riders.

Uber isn’t an outlier. It’s just early, large, and public. The same story is quietly playing out in midsize companies everywhere: budgets set with confidence in January, blown through by spring, executives staring at a usage bill they can’t tie to a single business outcome. If you’re trying to budget for AI right now, this is a primer on what actually matters, and what doesn’t.

First, the mistake almost everyone makes

It usually starts with fear, and the fear is real. You’ve been told about “shadow AI,” employees pasting company data into personal accounts with no oversight. And it is already widespread: in most companies, people are using personal AI tools for work faster than the company can sanction them. The work is already happening. It’s just happening off the books.

So the natural reaction is to get ahead of it: buy the enterprise licenses, flip them on across the company, hoping people will migrate over and slowly start adopting and driving useful things along the way. Then comes the second worry: token costs. The recent and more frequent stories emerging start to spook everyone. But here is the part the media gets wrong. Controlling token spend is not the hard problem. Setting limits, caps, and alerts is fairly straightforward. If runaway tokens are your biggest fear, you’re worried about the cheap problem.

> *The token bill is the symptom. The real problem is a missing strategy, management, and adoption plan.*

Because while everyone fixates on usage caps, they miss the bigger punchline: what the small group of companies actually winning with AI are doing to get a return that dwarfs everyone else’s.

The BIG perspective shift: a hiring decision, not a software purchase

Here is the mental model that changes everything. Imagine you hired five hundred brilliant new college graduates, genuinely smart, fast, capable people who happen to know nothing about your company, your customers, your data, or how you do things. Now imagine your entire plan for them was: bring them in tomorrow, turn them loose, and hope the rest of the company figures out how to use them well and cost-effectively. You would never do that. It's absurd. You'd have job descriptions, onboarding, managers, goals, reviews, and someone accountable for what they produce.

> *Budgeting for AI is not a software implementation. It is a workforce expansion. Missing that is the single most expensive mistake an executive can make.*

Software you install, configure, and forget. A workforce you recruit, onboard, manage, and develop. The companies blowing their budgets are buying a workforce and managing it like a software license. That is the whole problem in one sentence. So budget for it the way you'd budget for people. That means three things.

1. Budget for the management, not just the tool

The license and the tokens are the visible cost. The work that makes them pay off, the strategy, governance, policy, and the management framework that decides what AI does and who is accountable for it, is the invisible one. Early on, that invisible cost is often as large as the software itself. Companies that budget only for access are budgeting for the laptop and forgetting the employee.

2. Roll out by cohorts, against a roadmap with tracked ROI

You don't hire a 500-person team and skip the goals and performance reviews. Don't do it with AI either. The winners don't switch everything on at once. They sequence the rollout, starting with the executive team, super users, then the targeted teams solving the highest-value problems first, and extend company-wide from there. Every cohort is tied to a roadmap with clear ROI targets it measures against. Adoption becomes managed, not merely hoped for.

3. Make the objective a measurable ROI goal, not lower token usage

This is the counterintuitive one. Your objective is not to minimize spend. It is to drive the highest possible return in productivity and impact. Sometimes that requires heavy token usage; sometimes it doesn't. You cannot know which without a strategy and specific use cases, which is exactly why the roadmap matters more than the cap. A company optimizing for low usage is optimizing to get nothing done cheaply.

The three buckets of a real AI budget

A first-year AI budget should split across three categories, not pour everything into licenses. If nearly all of the money goes to access, you have funded usage but not value.

Budget bucket	What it funds	Why it matters
Access (~60%)	Licenses, model usage, enterprise accounts, admin controls	Moves people out of shadow AI
Deployment (~30%)	Use-case roadmap, data boundaries, workflows, integrations, cohort onboarding	Turns AI into a business capability
Management (~10%)	Governance, training, ROI tracking, feedback loops, evaluations, continuous improvement	Turns usage into measurable performance

What this looks like: wrong versus right

The wrong way	The right way
• Roll out company-wide overnight • Hundreds of users poking at slides and spreadsheets in scattered personal projects • Executives watch a cost-and-usage dashboard with no measure of output • No policy, nothing shaped to the business • Result: real spend, no provable value	• Budget set against an ROI-driven roadmap • Concrete 12-month targets by workflow (a 3x to 5x internal hurdle where the case is clear) • Rollout in cohorts: exec team, super users, roadmap teams, then company-wide • Cost and usage control plus real revenue and productivity indicators • Result: a workforce managed toward a number

> *One company is watching a meter. The other is managing a workforce toward a number. Same tools. Opposite outcomes.*

Measure it like a workforce, not like software

You would never judge a new team by how many badges were issued or how many emails they sent. Yet that is how most companies measure AI: seats activated, prompts submitted, tokens consumed. Those tell you the tool is on, not that it is working. Trade each one for a metric a manager would actually care about.

Don’t manage to this	Manage to this
Seats activated	Priority roles enabled
Prompts submitted	Workflows improved
Token usage	Cost per business outcome
Chat volume	Hours saved and cycle time reduced
Tool adoption	Department-level productivity gains
Experiment count	Use cases scaled or killed
User satisfaction	Revenue, margin, quality, risk, and speed impact

> *Low usage is not success. High usage is not success. Business impact is success.*

How much should you actually budget for AI?

All of that raises the practical question this article exists to answer: how much should you actually spend? Start with revenue, not seat count. For companies serious about AI, the real benchmark is not a handful of software licenses; it is a meaningful business-capability budget. PwC’s 2026 research puts numbers around it: AI leaders invest roughly 2.5 times as much of their revenue in AI as everyone else, leaders in software, banking, and media spend close to 5% of revenue, and outcomes begin to strengthen once investment crosses about 1.6% of revenue. Read 1.6% as a directional guide for committed investment and the 5% range as where ambitious leaders are heading, not as universal targets. The lesson is not to chase one magic percentage; it is that AI has become a real capability budget, not a software line item squeezed in as an afterthought.

What does that translate to in real dollars? It depends on maturity and goals, not just size. For a company around $100M in revenue, three rough tiers make a useful planning guide:

Program level	Annual budget (~$100M revenue)	What it funds
Starter / readiness	$250K–$500K	Enterprise AI setup, policy, training, first use cases, and basic management
Serious managed rollout	$750K–$1.5M	Company-wide access, cohort rollout, workflow redesign, governance, integrations, and ROI tracking
Aggressive transformation	$2M–$5M+	Redesigning core workflows, building AI agents, cross-system integration, and making AI a major operating advantage

For that same $100M company, this does not mean blindly writing a $1.5M AI check on January 1. It means recognizing that a serious AI program is likely a six- or seven-figure capability budget, not a $30-per-seat software purchase.

Just as important as how much is how you split it. PwC found the average enterprise still pours about 62% of its AI spend into technology, leaving roughly a third for process and change and almost nothing for training. That is the trap. If nearly all of your budget is licenses and tokens, you are starving the work that actually produces a return. A healthier steady-state allocation is closer to 60% access and technology, 30% deployment and workflow redesign, and 10% management, training, and measurement: the three buckets above, funded in proportion to the value they create. That split is the steady state. The first year, the one that builds the foundation, usually tilts the other way. Turning on an enterprise AI workspace is cheap and nearly instant; the heavy, one-time lift is everything that makes it usable and safe: the policies, the context and standards, the governance, and the cohort-by-cohort training of your people. Early on, deployment and management together can rival or outweigh the technology line, then settle toward the steady-state mix as usage scales and the setup is paid off.

> *AI budgets are not failing because companies are spending too much on AI. They are failing because companies are spending too much on unmanaged AI.*

Expect these figures to run much higher inside the functions AI is meant to transform. Gartner's 2026 CMO survey found marketing leaders allocating 15.3% of their budgets to AI, rising to 21.3% at the most AI-ready organizations, even as only about 30% felt ready to scale it. The point is not to copy 15% everywhere; it is that a department genuinely being reshaped by AI should expect to commit a real share of its budget to it, readiness and management included.

Again, benchmarks at this stage of market maturity are only a directional guide to help you make the best decision for your business. If you take nothing else from this article, take this: you need to budget

for a new workforce, not just license seats and tokens. Do that, and you are already ahead of many. If you haven't even thought about budgeting yet, you are not alone either, and we hope this helps you avoid some of the common mistakes we see with our customers. As always, for any questions or additional support, reach out to us at team@manageai.io. We are happy to share more knowledge and tools, such as budgeting worksheets, with no commitment.

Article Key Takeaways

1. Budget for AI as a workforce expansion, not a software purchase. The license is the cheapest part.
2. Split the first-year budget across access, deployment, and management. Funding only access funds usage, not value.
3. Token caps are the easy problem. The costly mistake is spending with no strategy behind it.
4. Roll out in cohorts against an ROI roadmap, the way you would onboard and manage new hires.
5. Measure business impact, not seats or tokens. In 2026 most companies still cannot tie AI to profit, and the ones who can are the ones who manage it.

Sources: PwC 2026 AI Performance Study, April 2026 (AI leaders invest about 2.5x more of revenue than peers, and around 5% of revenue in leading sectors); PwC AI investment benchmarking, 2026 (1.6%-of-revenue tipping point; 62/34/4 split across technology, process and change, and training); Gartner 2026 CMO Spend Survey, May 2026 (15.3% of marketing budgets allocated to AI, 21.3% at the most AI-ready organizations); Fortune and TechCrunch reporting on Uber's 2026 AI budget, 2026. Dollar ranges are illustrative planning guides for a roughly $100M-revenue company, not guarantees.

About ManageAI

Wave three is what we built ManageAI to solve for midsize growth businesses. Not another AI SaaS tool. Not an AI consulting firm. ManageAI is the first AI Workforce Management company, purpose-built to help businesses deploy, own, and manage AI workforces on infrastructure they control. We help you fund the management, sequence the rollout, and measure the return, so your AI budget buys outcomes instead of usage.

Ready to build a budget that buys outcomes? Learn more at manageai.io →