1Password

GUIDE

From blind spot to budget control

A practical guide for AI spend management across IT, Finance, and AI program leadership

Introduction

Organizations are adopting AI in pursuit of tangible outcomes: faster development cycles, competitive differentiation, and productivity gains.

But tying those outcomes back to AI spending and token use is often murky. IT can't trace the usage to specific tools and Finance has no breakdown by team or project. AI program leaders, who are accountable for AI strategy and outcomes, don’t know what the use cases running are costing. The question is simple: how much are we spending on AI, and are we getting the business value we expected?

Traditional SaaS spend is predictable. You count seats and multiply by price to understand the cost. And you have an opportunity to adjust the contract pricing at renewal. AI pricing works differently because it is based on usage. The cost varies by AI model, prompt complexity, context window, and whether the work is done by a person or an agent running autonomously.

Two teams doing similar work can generate costs that differ by an order of magnitude based on choices their tools make on their behalf.

Why AI spend behaves differently

Traditional SaaS
Predictable by design

- Priced per seat
- Renews on a contract
- Flat and forecasteble

AI Consumption
Unpredictable by nature

- Priced per token
- Bought everywhere
- Spiky and after-the-fact

Several dynamics make AI spend particularly difficult to manage.

It compounds without warning. An AI vendor can quietly shift a default model to a more expensive tier mid-cycle. An agent running an overnight task can burn through a week's worth of tokens on a loop it couldn't resolve. A team ramping a new use case can double a monthly burn rate before anyone thinks to check. Unlike a SaaS seat, which is a fixed cost until renewal, token spend accumulates in real time.

It's fragmented across vendors. Every AI vendor has its own billing model, its own metrics, and its own dashboard. Getting a credible total means logging into each portal separately, pulling exports in incompatible formats, and reconciling data that was never designed to be compared. By the time that reconciliation is done, it's already out of date.

Accountability lives in the gap between functions. IT sees the SaaS landscape and Finance sees invoices, usually after the money is already committed. AI program leaders understand the use cases but not the costs. No single function has a complete picture, and without one, no single function can act on it.

Accountability lives in the gap between functions. IT sees the SaaS landscape and Finance sees invoices, usually after the money is already committed. AI program leaders understand the use cases but not the costs. No single function has a complete picture, and without one, no single function can act on it.

Multiple blind spots, one shared challenge

Which tools are actually running. More than a quarter of knowledge workers use AI tools their employer didn't approve. Some reach the corporate bill through shared API keys, expensed subscriptions, or paid workspaces that never went through procurement. Others are simply invisible: not connected to the identity provider, not tracked in finance systems, not appearing in any managed inventory. Access may not be revoked when the person responsible leaves.

27% of employees have worked on AI-based applications that their employer did not approve.

What's driving the cost. Most AI vendor invoices arrive as a total. They show what was spent, not who spent it, on which model, or for what purpose. That gap between invoice and consumption makes it impossible to evaluate whether spend is justified, which project it belongs to, or whether the same output could have been achieved at lower cost.

Whether a spike signals a problem or reflects progress. When consumption jumps, it might reflect a high-value engineering sprint, or it might reflect an agent stuck in a retry loop, burning through tokens with no useful output. Without attribution data connecting spend to teams, models, and use cases, there's no way to tell the difference.

These aren't separate problems for separate functions to solve independently. They're shared blind spots with a shared cost: AI spend that can't be explained, governed, or optimized.

AI spend visibility is split across three teams

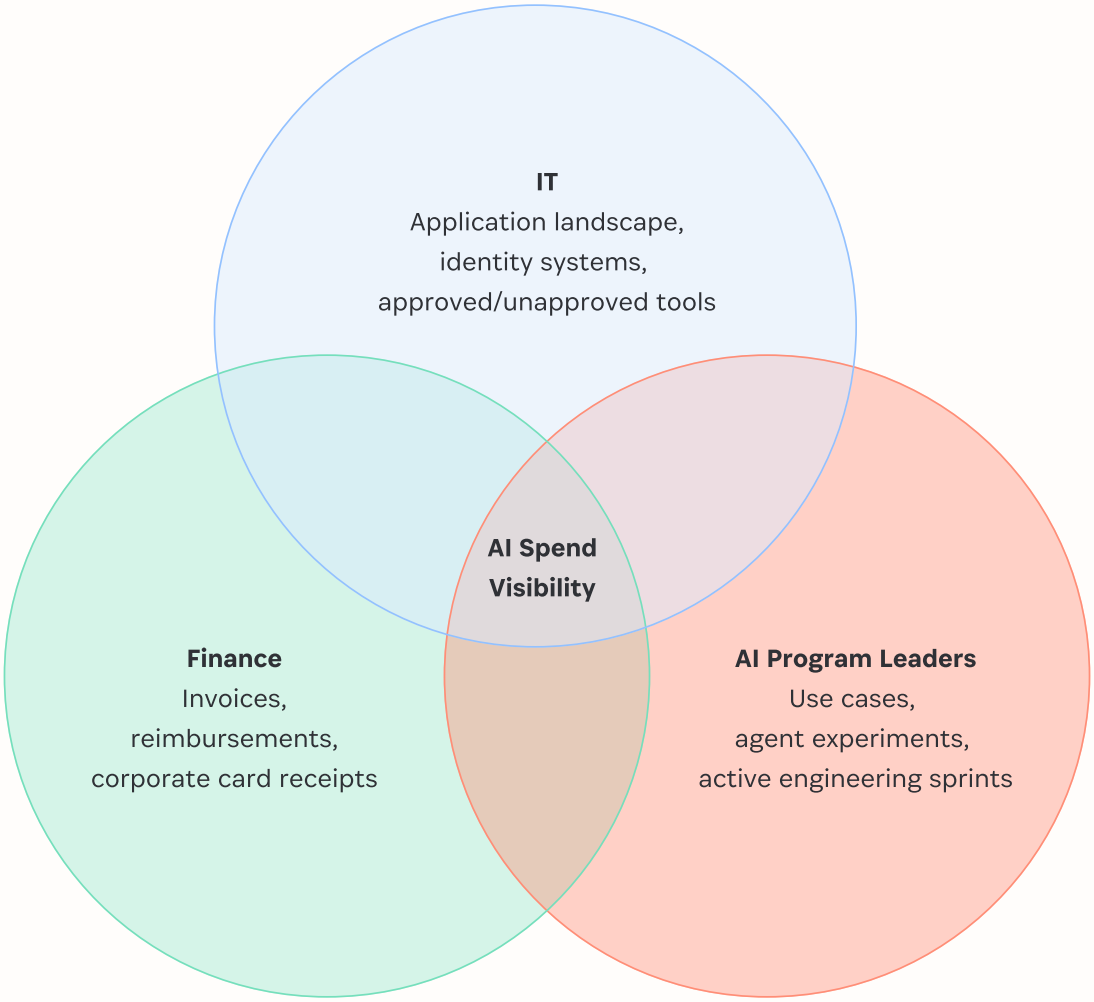

Implement guardrails grounded in how work happens

Without visibility into AI consumption and spend, companies often default to limiting access, blocking tools, and pulling back on experimentation. But that approach creates a new governance problem, by pushing AI usage onto unsanctioned shadow AI. The goal is to build controls that reflect how teams actually work, so the governed path is also the easy one.

Match each task to the right model. Model selection has the biggest effect on AI cost. Start with common use cases and test them on the least expensive model that can handle the work. Move to a more capable model only when the cheaper option misses the quality bar, then make those choices the default and define when exceptions make sense. Review the results regularly and adjust as needs change.

Agentic workflows need specific attention. Set per-run budgets and timeouts before deploying an agent. Then alert when it exceeds those limits or behaves unexpectedly. Track agent-generated usage separately from human usage so teams can spot runaway workflows quickly. Start with conservative limits and review real runs. Expand those limits only when the workflow consistently delivers the intended result.

Start with a baseline and then refine it. Thresholds and alerts should reflect actual patterns, but the first baseline is always a hypothesis. Set an initial ceiling based on acceptable cost, risk, and intended workloads. Then monitor usage and refine it as legitimate patterns emerge. If current consumption is inflated by waste or experimentation, don't let the highest observed usage define the baseline. Use the data to distinguish necessary work from excess, and update the controls accordingly.

1Password makes these controls concrete. Finance and IT can set budgets per vendor, configure alert thresholds, and track daily burn rate against them, with depletion estimates that flag a prepaid balance heading for early exhaustion before the contract term ends.

From reactive to proactive

When AI spend is governed with visibility, control, and attribution, the questions leadership is already asking become answerable. The full inventory of AI tools in use, including those that never went through procurement, stays current rather than surfacing during an incident review. Spend breaks down by team, model, and project before the invoice arrives, not after. And the question of whether AI investment is producing proportional value has an answer backed by attribution data, not assumptions or guesses.

Governance without a clear sense of what AI investment is supposed to produce tends to default to restriction: caps, limits, and cutbacks applied without context. Tying token consumption to business outcomes is what gives governance its direction. When teams can articulate what a given AI investment is driving toward, Finance can evaluate spend against something real, not just a number. That foundation is what separates organizations that govern AI spend from ones that simply watch it.

When IT, Finance, and AI program leaders are working from the same data, the gap between what's being spent and what's understood closes. Spend that was once invisible becomes attributable and budgets move from reactive to forecastable. Model choices that were once defaults become deliberate decisions. When the invoice arrives or the questions are asked, Finance, IT, and AI program leaders have the confidence to answer.

How 1Password controls AI spend

1Password SaaS Manager addresses these challenges through AI Spend and Consumption Management. This capability gives IT, Finance, and AI program leaders a single, normalized view of what AI is costing, who is driving it, and when budget risk is approaching. Built on the same foundation used to govern the broader SaaS portfolio, it normalizes usage and spend data across AI vendors into one view, refreshed daily.

Visibility. A single dashboard shows AI consumption and spend across vendors such as Cursor, Anthropic (Claude), and OpenAI (ChatGPT), and Amazon Bedrock. That includes tools that entered the environment through expense reports, browser extensions, or personal accounts, not just those provisioned through formal procurement. The dashboard organizes spend by vendor, team, user, and model, giving IT a complete picture of the AI estate rather than a partial view limited to sanctioned tools. For organizations managing multiple AI vendors simultaneously, the data is normalized into one view that refreshes daily, so the picture is current rather than a reconciliation of last week's CSV exports. IT, Finance, and AI program leaders work from the same data rather than separate, partial views.

Control. Finance and IT can set budgets per vendor, configure alert thresholds by percentage, and track daily burn rates against them, eliminating surprises. For any prepaid token balance, the platform calculates an estimated exhaustion date at the current burn rate. Over-budget, depletion-risk, and data-gap notifications arrive automatically in Slack and email, so no one has to remember to check.

Attribution. AI spend breaks down by vendor, team, user, API key, and model, making it possible to see exactly where consumption is concentrated and where optimization opportunities exist. That includes identifying when teams are defaulting to the newest, most expensive model for work that a less costly one could handle. Attribution by model and user has a dual purpose: it lets organizations direct more capacity toward the work producing real value and rein in the usage that isn't.

Watch the demo to learn more

What sets 1Password apart

The tools most organizations already have weren't designed to answer the questions AI spend is now raising.

- Expense management platforms see AI costs when they hit a corporate card, not before.
- Procurement tools track contract values but have no line of sight into the token consumption that determines whether that contract will be enough.
- Vendor dashboards exist, but each one measures differently, exports in its own format, and goes stale the moment you close the tab.

What's missing is a tool that connects AI consumption to the rest of the software governance picture: managing discovery and access alongside spend. 1Password SaaS Manager provides that layer. Since it already governs the broader SaaS portfolio, AI tools discovered in the environment don't live in a separate system. They surface in the same platform, matched against a library of 40,000+ pre-populated application profiles, with risk and compliance context attached. Organizations that have connected SaaS Manager to their identity provider and finance systems have discovered thousands of applications they had no prior visibility into, often in minutes.

The result is a governance framework that goes beyond spend. Organizations know not just what AI is costing, but which tools are active, who has access, and whether that access has been governed. That combination is what makes AI spend manageable rather than just visible.

Conclusion

1Password

AI spend is no longer a line item Finance can manage reactively. The combination of consumption-based pricing, fragmented vendor data, and accelerating adoption has created a visibility gap that individual vendor dashboards and manual reconciliation can't close.

1Password SaaS Manager gives IT, Finance, and AI program leaders the shared foundation they need: a normalized view of AI consumption across vendors, budget controls that act before the invoice arrives, and attribution data that connects spend to the teams, models, and use cases driving it. AI Spend and Consumption Management is built on the same platform already managing the broader SaaS portfolio so there is no new vendor to evaluate, no custom deployment to manage, and no additional cost for existing SaaS Manager customers.

The organizations that establish this foundation now will be demonstrating business outcomes, not explaining overruns.

To see how the AI Spend and Consumption Management in SaaS Manager can help you get AI spend under control, contact our team to schedule a demo.