harness

2026 State of AI in FinOps

AI Spend: Rising Fast & Running Blind

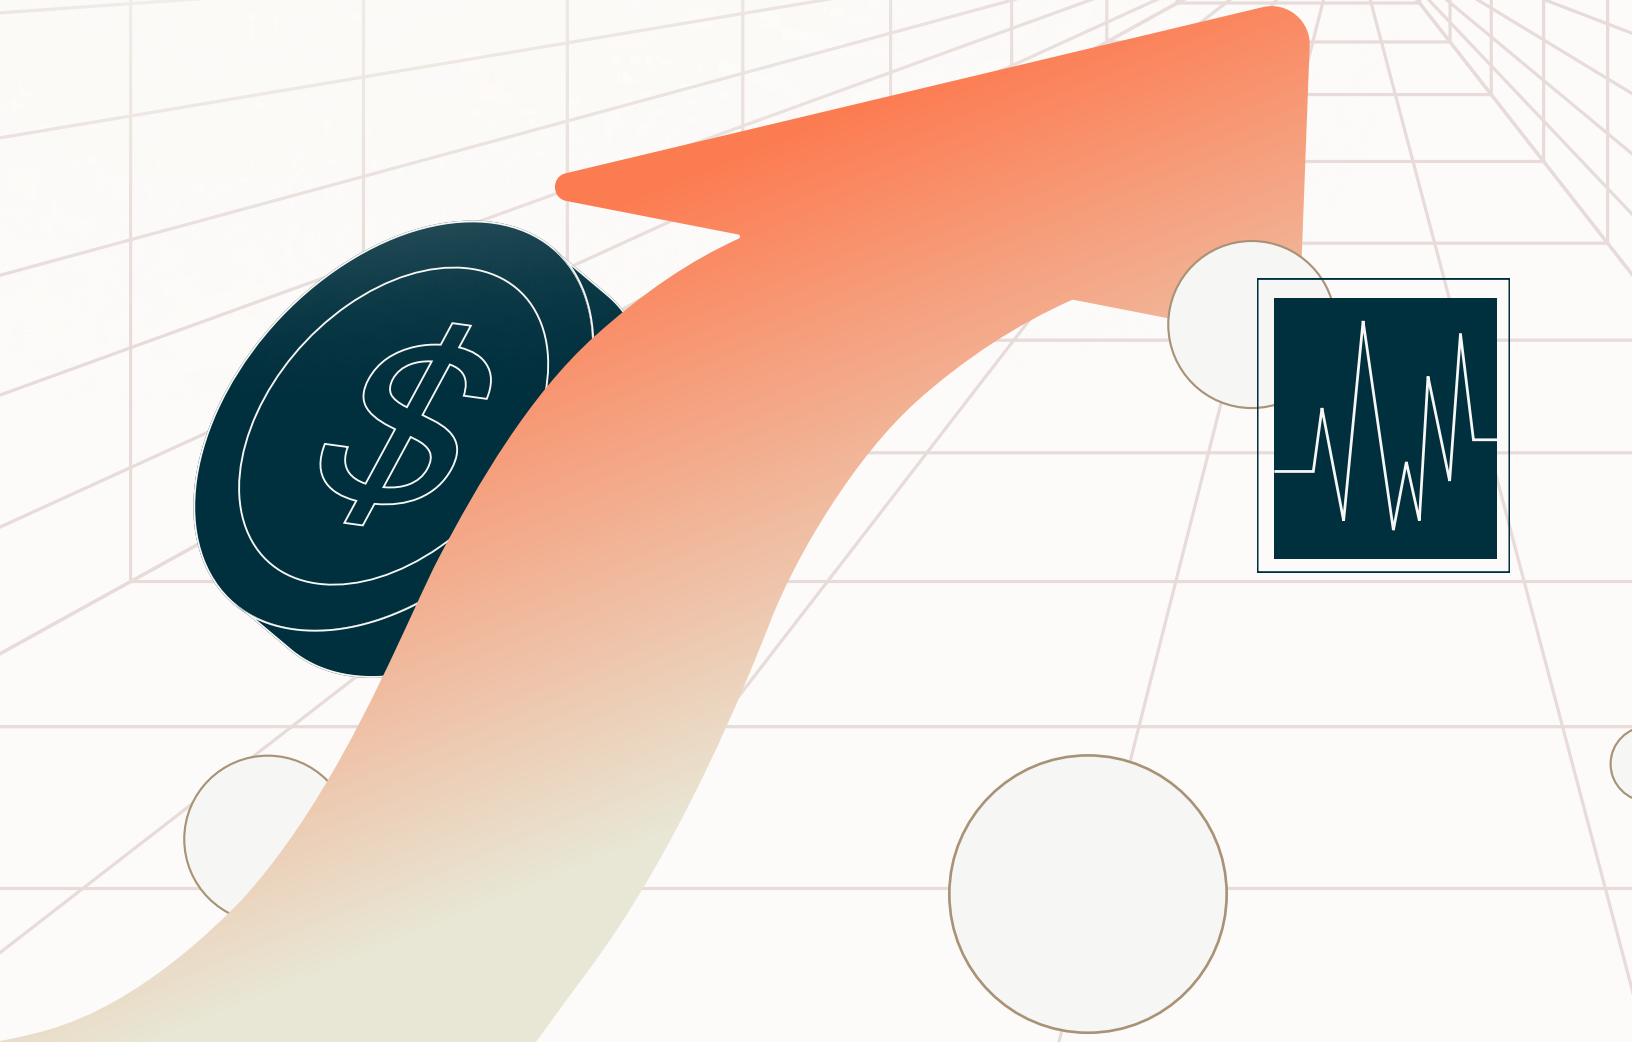

Executive Summary

AI isn't just being adopted anymore. It's now a $1M-a-month line item for 20% of enterprises, and two in three organizations spend more than $250,000 a month. At that scale, AI isn't a tool cost; it's a capital decision that requires the same governance, attribution, and ownership infrastructure that took cloud computing a decade to build. Today's businesses, however, can't wait that long.

We surveyed 700 engineering leaders and practitioners across five countries this spring to ask about their organization's FinOps practices. Specifically, we wanted to understand: where is AI spend actually being attributed, and who owns it? What we found isn't an emerging cost concern, it's an established budget reality still being treated like an experiment. Spend is climbing fast and broadly, across every category and provider at once, while the ownership, visibility, and governance structures needed to manage it haven't caught up.

The result is a quietly expensive blind spot: organizations estimate that roughly a quarter of all AI spend is wasted. Even worse: when a cost spike hits, most teams can't trace its source in under a day, and nearly a third need a full week to determine the issue. Common to both problems is the same root cause of fragmented ownership. Four different functions each hold a piece of the bill, and no single function owns the total.

If you're in a FinOps, platform, or engineering leadership role watching AI costs show up as a line item nobody can fully explain, this report is both a benchmark for where the industry stands and a practical roadmap for what to fix first.

Key Findings

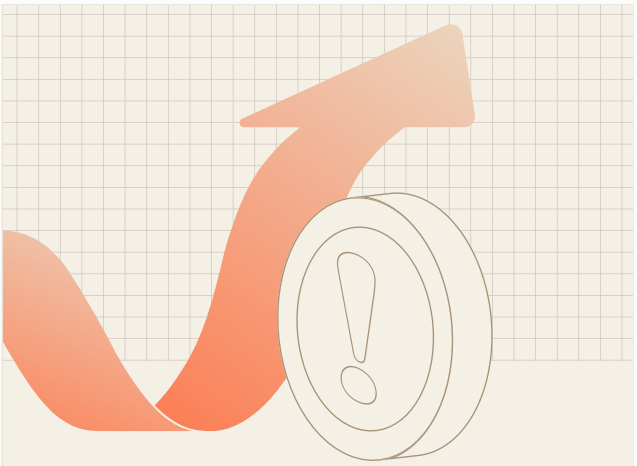

72%

hit a surprise AI bill in the past year.
1 in 3 have been caught off guard multiple times.

52%

say there's no clear AI cost owner.
Accountability is split four ways across engineering, FinOps, finance, and IT.

26%

of all AI spend is wasted.
Consistent across every category: infrastructure, software, models, and services.

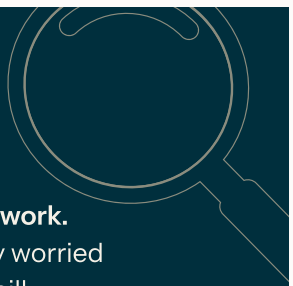

56%

say AI budgeting is guesswork.
47% say they're constantly worried about their next AI token bill.

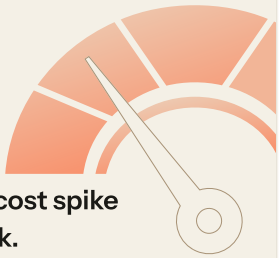

79%

need a full day+ to trace a cost spike
32% would need a full week.
At $1M/month, that delay is expensive.

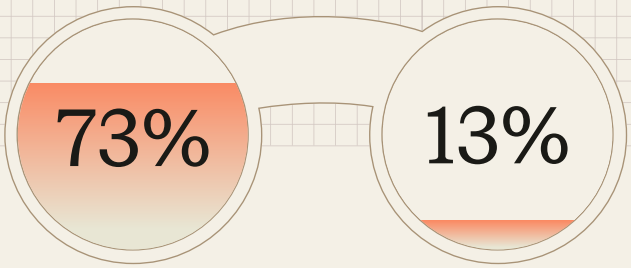

You can't enforce a policy when you can't see the data. The gap between 73% and 13% is the enforcement problem.

Have AI cost policies

Have basic visibility into cost

PART 1

The AI Spend Reality

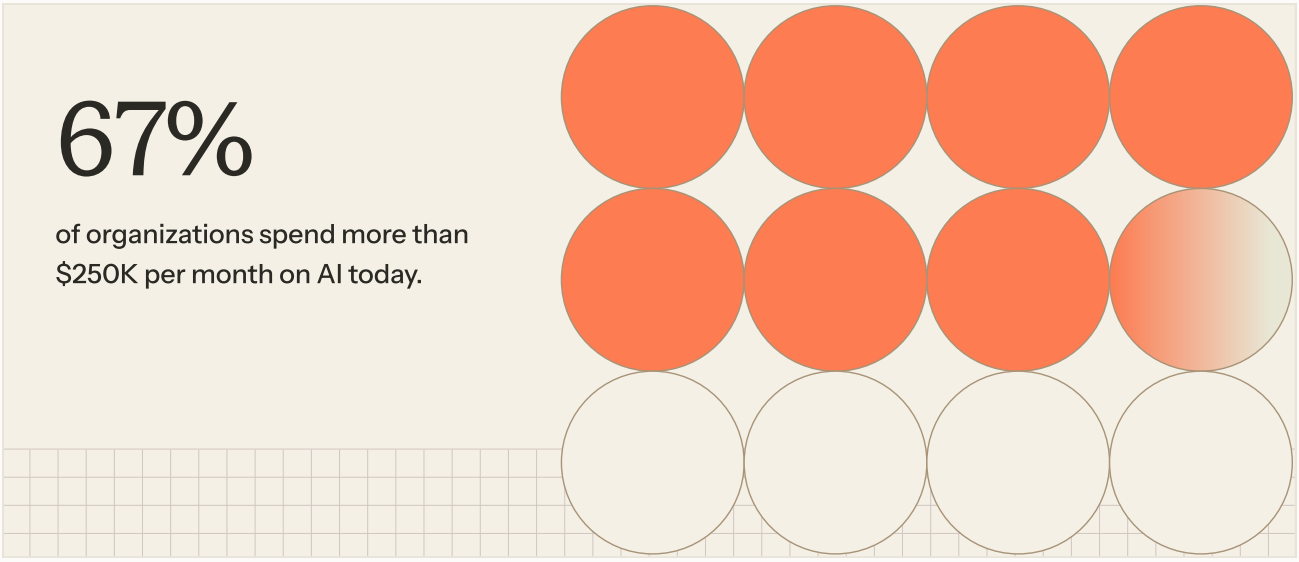

The spend is here, it's substantial, and it's accelerating across every category at once: infrastructure, services, software, models, all rising together. This is backed by the data: 80% of respondents say overall AI spend increased in the past six months, and 78% expect it to keep climbing.

What makes this moment different from past technology spend cycles is the breadth. When cost grows simultaneously across infrastructure, software, and models, AI stops behaving like a point solution and starts behaving like a platform-level investment, the kind that needs real ownership structures and attribution built underneath it. Today, that infrastructure is exactly what most organizations are missing.

Areas of AI Spend Increases Over the Last Six Months

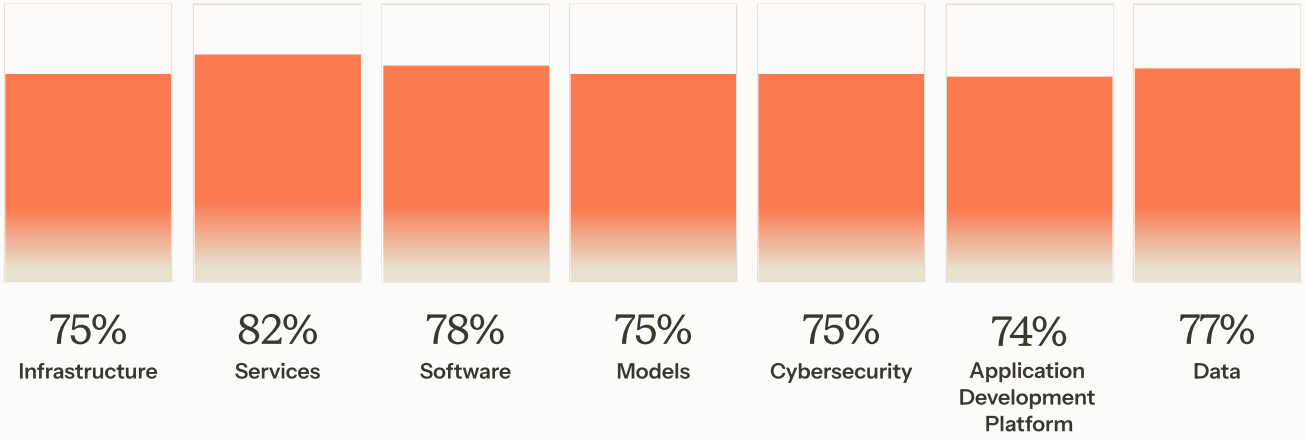

Part of the reason is that **AI spend isn't really one thing to manage, it's several.** Most enterprises run three or more major AI providers simultaneously, each with its own billing model, making a single trustworthy cost view structurally hard to build.

Provider Usage In Production

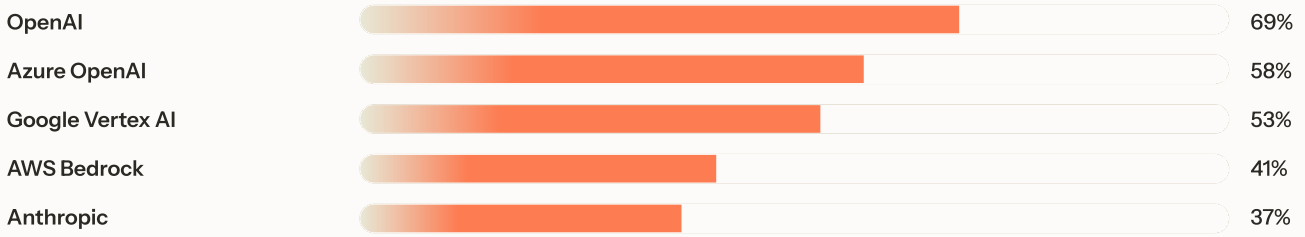

The typical organization isn't choosing between providers — it's managing all of them at once.

That overlap is the real culprit: the bill arrives in pieces, from different sources, in different formats, on different terms.

Visibility issues persist beyond provider sprawl. The top two cost drivers are productivity tools — packaged AI/SaaS solutions like Copilot and Glean, and coding assistants — not the customer-facing AI features that often get blamed. These costs look like software licenses, not infrastructure, and traditional FinOps processes weren't built to track them.

Approximately 3 in 10 organizations (32%) already attribute more than a quarter of their total cloud spend to AI, which is enough to make this a CFO-level conversation rather than a back-office one.

This level of growth would be manageable if the infrastructure to track it had kept pace. It hasn't, and the next section gets at why – not because nobody is paying attention, but because no one knows who is responsible for watching.

PART 2

The Ownership Crisis

If you ask five people in a typical enterprise who owns AI costs, you'll get five different answers. The data confirms that impression, and shows the confusion is structural, not incidental. More than half of respondents (52%) say there's no clear, dedicated AI cost owner in their organization.

But the deeper problem isn't absence, it's diffusion. Four different functions each hold a piece of the bill, with no single owner accountable for the whole.

AI Cost Accountability

When asked who is accountable for managing AI costs, no single function came close to a majority.

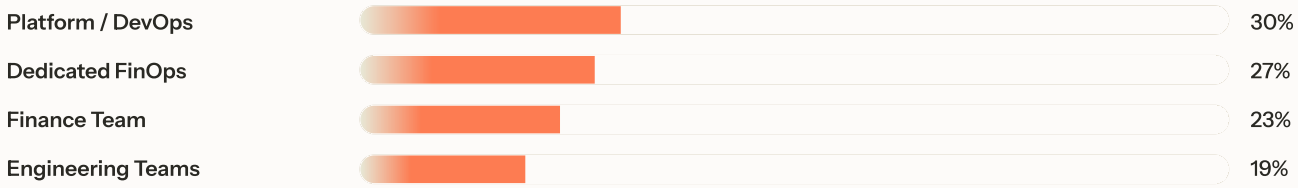

1 in 2

agree there's no clear, dedicated AI cost owner in their organization.

35%

Engineering and platform teams hold the most influence over AI spending decisions. Influence over spending ≠ accountability for the bill. That gap is where surprises are born.

When the bill arrives, the first conversation is about who should have caught it, not how to fix it.

There's a structural reason ownership is harder to assign for AI than it was for previous technologies. Cloud costs track to infrastructure, servers, storage, and network, so ownership follows the asset. AI costs track to usage, driven simultaneously by product decisions, engineering choices, model selections, and real-time user behavior. No single function owns all of those inputs, yet all of them land on the same invoice line.

That mismatch becomes sharper when you look at who actually drives the spending versus who's left paying it: engineering and platform teams have the most influence (35%) over AI spending decisions, yet platform/DevOps is accountable for 30% and FinOps for 27%. Influence and accountability aren't aligned, and when the people making cost-driving decisions don't feel the consequences of them, they are not incentivized to be responsible stewards of technology.

PART 3

Cost Surprises & Visibility Gaps

Cost Surprise Reality

If your AI spend doubled overnight, how long would it take to identify the source?

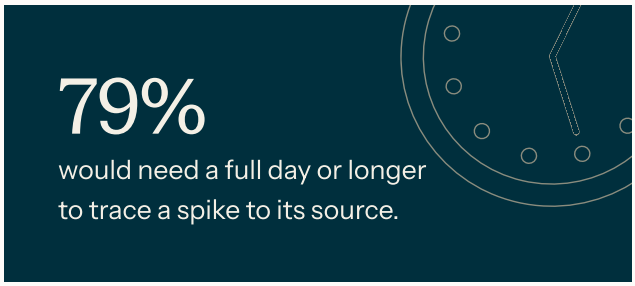

For 1 in 3, tracking down the cause takes a week or more — while the bill keeps running.

The ownership problem creates a visibility gap that makes cost surprises inevitable and diagnosis painfully slow. Nearly half of respondents (45%) say unexpected cost increases happen frequently or constantly. When surprise becomes the norm, it stops triggering action, and that desensitization is the real risk.

But visibility doesn't solve everything. Knowing a spike happened is only the first indication of a potential problem. Understanding the root cause in near-real-time is where most organizations fall short.

When A Spike Hits - How Long To Find The Cause?

If your AI spend doubled overnight, how long would it take to identify the source?

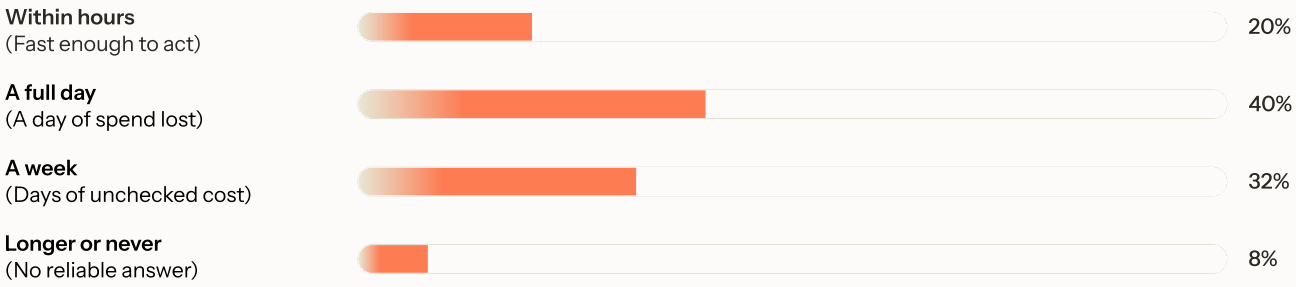

1 in 5
Diagnose in a day or less.

At **$1M** /month
a week of confusion about why costs spiked is a serious financial exposure.

A fifth resolve within hours. But 40%, the largest group, need a full day, nearly a third need a full week, and 8% never reliably find the cause at all. At $1M a month, a week of confusion is a serious financial exposure, not a minor delay.

The diagnosis failures cluster around the same gap: organizations can't tell what changed (whether in code, configuration, prompts, or models) fast enough to connect cause to effect. More than 4 in 10 respondents still manage AI costs with spreadsheets, a stopgap built for early spend levels that's still standing years later.

That gap between expectation and reality shows up most clearly in how organizations plan ahead. More than half (56%) say anticipating AI spend is based on guesswork, not data. This chronic uncertainty isn't contained to balance sheets: it affects hiring, roadmaps, and model selection. When the cost of a decision is unknown, most organizations default to guesswork rather than preparing with better information.

PART 4

Building Blind

For 1 in 3, tracking down the cause takes a week or more — while the bill keeps running.

Structural gaps in ownership and visibility explain a significant portion of the AI cost management problem. But there's a parallel, bottom-up dynamic that policy alone can't fix: the people actually building AI features often have no idea what those features cost to run, and the organizations they work in are actively incentivizing them to use more AI, not smarter AI.

More than half of respondents (57%) say their organization encourages "tokenmaxxing," where employees are pushed to maximize AI use regardless of tangible value. Maximizing usage made sense when organizations were trying to build a habit, but that adoption phase has expired and productivity has plateaued.

The Visibility Gap

26%

of overall AI spend is estimated to be wasted — consistently, across every category.

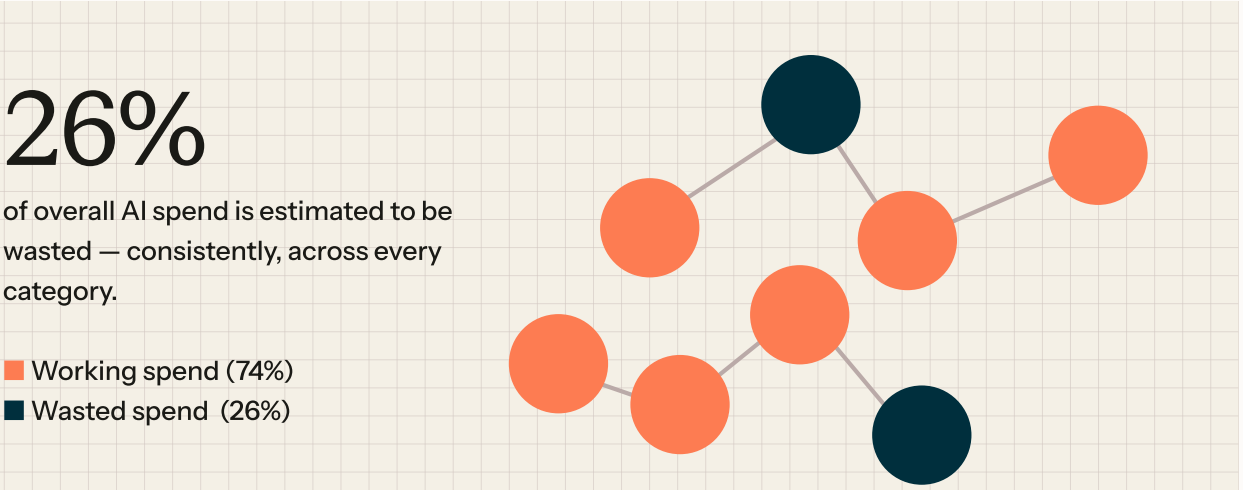

What AI Waste Costs Every Month

$500K monthly budget

$130K wasted

$1M monthly budget

$260K wasted

Unlike cloud waste, AI waste is harder to detect — it's embedded in usage patterns, not infrastructure configuration.

When users are incentived to maximize consumption without being provided the cost signal, waste and value become indistinguishable. The cost of that blind spot is concrete: organizations estimate that 26% of all AI spend is wasted, fairly evenly across every category, software, models, infrastructure, services alike. For an organization spending $1M/month — the reality for 1 in 5 respondents — that's $260,000 a month with no measurable return.

Engineers who don't factor cost into their decisions aren't doing it because they don't care, the data simply isn't in front of them when decisions get made.

Engineer Cost Awareness

56%

of engineering leaders say that teams don't have cost in mind when building AI features.

45%

average share of engineers who understand what they build costs.

The problem isn't that engineers don't care. It's that the cost data isn't visible at the moment decisions get made.

This isn't a motivation gap. In our 2025 FinOps in Focus report, 62% of developers said they wanted more control over cloud costs, and there's no reason to believe AI is any different. The problem is visibility, not intent: cost data isn't reaching engineers at the moment decisions actually get made. Prompt design, model selection, retry logic: none of those moments currently carry a cost signal.

The fix isn't more policy. It's making cost visible at the moment of decision.

PART 5

Governance, Maturity, & What's Next

AI Cost Governance

The Real Tension

You can't enforce a policy when you can't see the data. The gap between 73% and 13% is the enforcement problem.

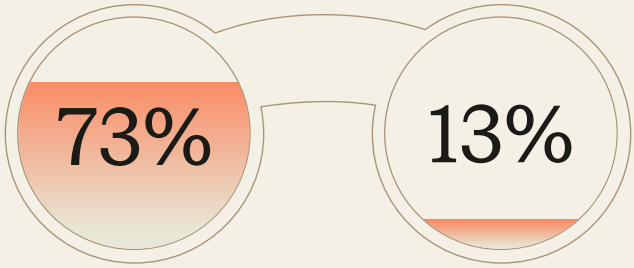

The data paints a consistent picture of an industry with the right instincts about AI cost management, but without the operational practice to back them up. Most organizations have governance on paper: 73% say they have AI cost policies, but only 47% fully enforce them. That's a 26-point gap between writing policy and following it. Policies that don't surface in daily tools get ignored, and the incentive to move fast usually wins.

4 in 5 organizations have not yet reached mature, company-wide AI cost management. Only 21% describe themselves as fully mature, and only 26% have a robust method for measuring the business value of AI spend at all.

AI Cost Management Maturity

The spend is growing faster than the governance structures designed to manage it.

21%

describe their AI cost management as fully mature and company-wide.

26%

have a robust method for measuring the business value of AI.

What's Next

The organizations that have reached full maturity didn't try to fix everything at once. They followed a recognizable sequence: name an owner before buying a tool, build a single cost view before optimizing, embed cost data into engineering workflows before writing policy, and establish unit economics tied to meaningful business outcomes before measuring ROI.

That sequencing matters because the alternative – reviewing costs occasionally – doesn't match how fast AI spend actually moves. 42% review AI costs only quarterly; for a category that can move materially in a week, that's the operational equivalent of flying blind. The most meaningful shift an organization can make does not always mean procuring a new tool: rather, treating AI cost management as a development-time discipline instead of a post-production afterthought. In other words, make the people who build or deploy AI features responsible for the operational cost of those features, and most critically, give them the data to act on it.

Methodology

About This Survey:
This report is based on an online survey of 700 engineering leaders and practitioners about their organization's FinOps practices, conducted in May to June 2026.

Company Size

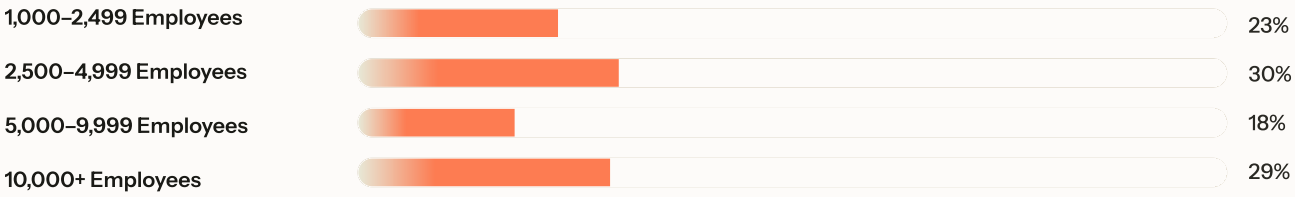

Annual Revenue

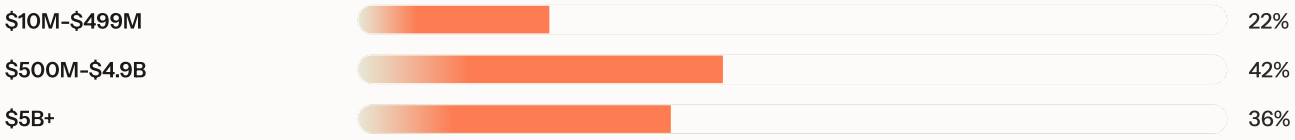

Countries

United States **n=300** | United Kingdom **n=100** | France **n=100**

Germany **n=100** | India **n=100**

These aren't startups experimenting with AI. Every organization in this survey has at least 1,000 employees and real, recurring AI spend.

This report is based on an online survey of 700 engineering leaders and practitioners about their organization's FinOps practices, conducted in May to June 2026. All respondents work at organizations that actively use AI/LLM services and are employed in software engineering/development, DevOps, IT operations, or executive leadership roles.

Sample Breakdown

Countries: United States (n=300), United Kingdom (n=100), France (n=100), Germany (n=100), India (n=100).

Company size: 1,000–2,499 employees (23%), 2,500–4,999 (30%), 5,000–9,999 (18%), 10,000+ (29%).

Annual revenue: $10M–$499M (22%), $500M–$4.9B (42%), $5B+ (36%).

Role type: Decision makers (50%), Managers (50%).

Primary function: IT operations/infrastructure (49%), Software engineering/platform delivery (39%), Executive leadership (13%).

Percentages may not sum to 100% due to rounding. Multi-select questions are noted where applicable. Where combined percentages are cited (e.g., "frequently/constantly"), this represents the sum of those response categories. Data marked with an asterisk (*) in the underlying dataset indicates a sample size too small for reliable reporting and has been excluded.

Note From the Authors

When we published the FinOps in Focus report last year, the market was focused on cloud costs: attribution, governance, getting engineers to care. That work isn't done. But something significant has shifted.

In the past twelve months, AI spend has moved from a line item that occasionally surprised people to a budget category that regularly does. The patterns, invoice shock, ownership confusion, governance gaps, are the same ones we saw in the cloud ten to fifteen years ago, compressed into a much shorter timeframe.

Patrick Brogan

What I keep coming back to is that the hardest part of this problem is not technical. It is cultural and organizational. The tooling is catching up. The harder work is getting engineering teams to build with cost in mind from the start, not as a constraint, but as a design principle. The organizations that solved this in the cloud didn't just buy better tools. They They shifted left and made financial data visible at the point where architectural decisions get made, not after the invoice arrives. That is the journey AI cost management needs to take.

Director, FinOps Advisory, Harness | FinOps Certified Practitioner, FinOps Foundation

---

Harish Doddala

A year ago, our customer conversations were about cloud cost attribution, commitment coverage, and rightsizing. Those still happen, but a more urgent question keeps interrupting them: why did my AI bill do that, and how do I stop it from happening again? The questions changed, so we ran this study.

What surprised me wasn't the size of the spend or the speed of the growth. It was how consistent the gaps are across every size and geography. The visibility problem, the ownership problem, and the forecasting problem show up whether you're spending $300K a month or $3M. Time won't close a gap like that. It's structural, and it takes deliberate investment in the right infrastructure to fix.

Use this report to find out where you stand, and where to start.

SVP & GM, Cloud & AI Cost Management, Harness