

AI Software Spending in 2026

NDR Slows, Pricing Resets,
and a Second Wave of Growth Follows



Executive Summary

The first sign of restraints on AI-native software spending hit in Q2 2026, and that's the best thing that's happened to buyers since the launch of generative AI.

Net dollar retention for AI-native vendors peaked at 136% in April, then fell for two straight months, and has leveled off in July. Wallet share going to AI tools kept climbing anyway, especially at the Enterprise level. Foundation model pricing structures flipped, with Anthropic setting the pace ahead of OpenAI, for the first time in the two companies' history. And established vendors in security, compliance, and financial planning are posting some of their fastest growth in years, not because they sell AI, but because they're the ones cleaning up after it.

Here's Tropic's read on what the first half of 2026 actually meant for the buyers footing the bill.

Key Takeaways

1

AI-native NDR fell two months straight in May and June, after peaking at 136% in April, the clearest sign yet that AI-native vendors are finally being evaluated like every other software purchase

2

Enterprise flipped from the smallest AI-native wallet share of any segment (1.4% in January 2025) to leading on the top segments of spend (5% in July 2026). Adoption plateaued; depth of use didn't.


3

Anthropic broke the pattern of following OpenAI's pricing moves for the first time. And both have made it loud and clear that the usage line is where the real spend lives.

4

Fastest-growing vendor data shows a second wave of growth behind AI: established security, compliance, and planning vendors accelerating because of AI's operational fallout, not because they're AI products themselves.

CHANGES IN MEDIAN NET DOLLAR RETENTION



Is AI-Native NDR Actually Slowing Down? Here's Why That's Not Bad News

Across Tropic's customer base, median net dollar retention for AI-native vendors (companies where AI is *the* product, not a feature) hit an all-time high of 136% in April 2026. Since then it fell two months running, 133% in May, 129% in June and stayed flat at 129% in July.

For context, 100% NDR means a vendor is holding flat. Anything above 110% signals real organic expansion. So, even at 129%, these companies are still growing customer spend faster than almost anything else in the software market. But the direction changed, and the reason why says more about buyers than it does about the vendors.

AI-native NDR just posted its first back-to-back decline in over a year, which shows the market might finally be growing up.

Adoption has plateaued but depth hasn't

The AI-native surge that drove NDR to 136% wasn't subtle. Enterprise AI-native spending grew 94% year over year through Q1 2026, while traditional SaaS growth cooled to 8%. On February 3, that gap triggered what the financial press started calling the "[SaaS apocalypse](#)," a 48-hour window in which roughly \$285 billion in software market cap evaporated as investors concluded that agentic AI would compress the number of seats companies needed to buy. AI-native vendors were the clear winners of that repricing, and our NDR data through April reflects it. Budgets moved toward them fast, with few of the usual friction points slowing expansion down.

That friction is back, kind of. Starting in late May, a wave of reporting, [led by an Axios story on May 28, described what's now being called "AI sticker shock"](#). Corporate leaders confronting ballooning token bills with uncertain productivity gains attached. A detail from that report noted one company spent half a billion dollars in a single month after failing to put usage limits on employee AI licenses. It was also reported [that Uber burned its entire AI budget for a year](#), in just 4 months. [Microsoft canceled most of its Claude Code licenses over cost](#). A [KPMG Global AI Pulse survey](#) released in late June, found that only 7% of leaders can point to established, measurable AI ROI. That gap is exactly what's pulling CFOs, not just CTOs, into reviewing every AI line item.

Where this shows up in the numbers

That shift explains the AI-native deceleration better than anything vendor-side. The April peak featured budgets expanding on enthusiasm. The June and July numbers are starting to show budgets expanding on evidence. CFOs, and everyone in-between, are finally asking AI-native vendors the same question they ask all technology companies: what did we actually get for this?


Meanwhile, the two other cohorts we track kept moving in their existing direction.

- Hybrid SaaS + AI companies, legacy vendors with credible AI capability built in, slipped from 101% in April to 99% in June and 98% in July. Still the "treading water" story we've told before, just drifting slightly negative for the first time this year.
- Primarily SaaS vendors, the ones without a real AI story, fell from 93% to 91% over the same stretch and are closing in on the point where NDR turns net negative.

The mix of sticker shock and unclear ROI explains why AI-Native spend still grew overall in Q2. Q3 will be the real test.

Compared to total changes in software and AI spend:

YOY CHANGE IN AVERAGE SPEND BY PRODUCT CATEGORY



Overall Year Over Year Change in Software and AI Spend




What's next for AI-Native suppliers?

The easy read of this data three months ago was that AI-native wins, legacy loses, simple as that. The Q2 and early Q3 numbers complicate that story slightly. **AI-native vendors aren't losing budget, but we are seeing them lose the complete blank check.** The same scrutiny that's been forcing legacy vendors to justify AI-driven price increases is now reaching the AI-native vendors themselves. This is a sign the category is finally being evaluated the way every other software purchase is – against a business case, not just a trend line.

For procurement, that's an opening. Look at who is providing real usage data, real outcomes, and real cost discipline. They will continue to see 129%+ retention, and are the ones you want to push your organization to keep working with.

OVERALL SHARE OF WALLET




How Much of the Software Budget Is Actually Going to AI-Native Vendors?

As of July 2026, the average share of software budget going to AI-native vendors sits at 4.8%, more than double the 2.3% it was in January 2025. The top decile of spenders, the outlier group, crossed 10% of wallet for the first time in June, and hit almost 13% in July 2026, up from 5.3% eighteen months ago. That confirms what we've been seeing.

AI-native spend is still a small slice of the average company's stack, even as the growth rate stays remarkable.

What's new is where that growth is actually coming from. Not new adopters, but existing ones spending deeper into tools they already use.

GROWTH IN PERCENTAGE OF ORGS SPENDING ON AI-NATIVE TOOLS



Adoption has plateaued but depth hasn't

As of July 2026, 67% of organizations are spending money on AI-native tools, up from just 12% in January 2023 when AI started to take-off. That climb was steep and steady for two and a half years, adding roughly two points of adoption a month through most of 2024. It's since flattened out. Adoption has hovered in a tight band between 64% and 67% since January.

That plateau matters for how we read the wallet share numbers. The average share of wallet is still climbing even though the population of adopters mostly isn't. In other words, more companies aren't discovering AI-native tools, which makes sense. The same roughly two-thirds of companies that already adopted AI, is spending meaningfully more within their existing footprint. Adoption found its ceiling. Intensity of use didn't.

64–67%

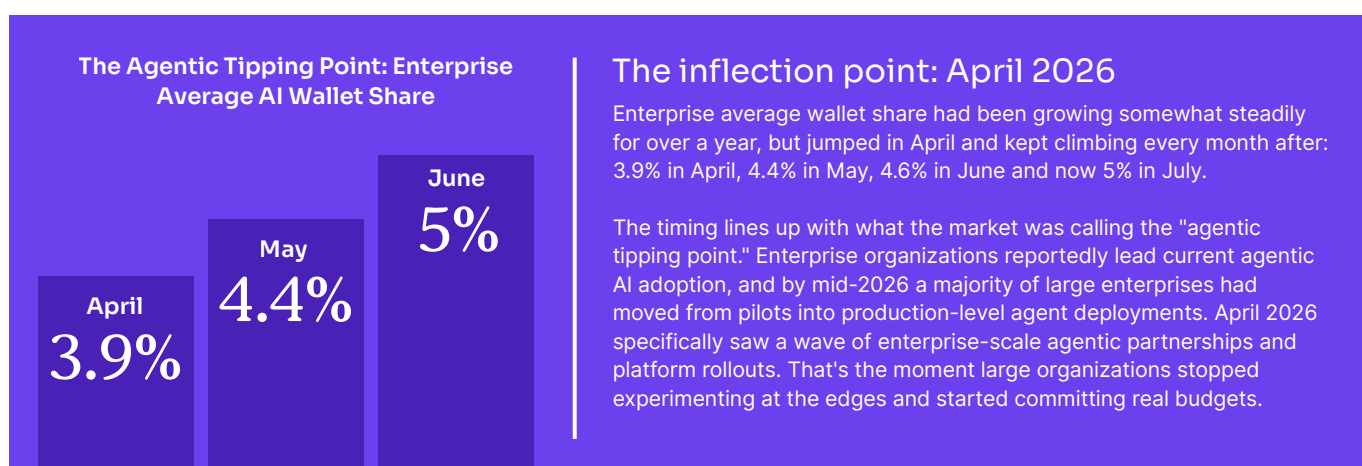
range of adoption percentage
among organizations since January 2026,
indicating an adoption plateau.

Which Company Size Is Spending the Most on AI-Native Tools in 2026?

Enterprise has gone from laggard to leader

Breaking wallet share out by company size shows a pattern that held for most of 2025 close to flipping entirely in 2026. In January 2025, Enterprise companies had, by a wide margin, the smallest AI-native wallet share of any segment: 1.4% average, compared to 3.9% for SMB and 4.1% for Midmarket. Enterprise was the laggard. SMB and Midmarket, quicker to adopt turnkey AI tools without the procurement friction that comes with large organizations, were out ahead.

By July 2026, Enterprise average wallet share reached 5%, even with SMB (5.0%) and slightly ahead of Midmarket (4.4%). On P90, the measure that captures the most aggressive spenders, Enterprise has already taken the lead outright: 13.7%, versus 13.3% for SMB and 11.7% for Mid-Market. Enterprise hasn't caught up across the board yet, but at the outlier end, it's now ahead of the pack.



What does this actually mean?


Putting this all together, two-thirds of organizations that were going to start spending on AI-native tools already have. The next phase is about who deepens their spend, and right now Enterprise is deepening the fastest.

This doesn't mean Enterprises have 'figured out' AI ROI any better than anyone else. CFOs are reviewing every AI line item across every segment, and Enterprise is no exception. **What it means is that when Enterprise commits, it commits at a scale that reshapes the aggregate numbers almost by itself.** Watch this segment closely over the next two quarters. If Enterprise NDR and wallet share both keep climbing together, that's going to make a huge difference in what suppliers survive this complex environment. If wallet share keeps climbing while NDR stalls, that proves the sticker-shock story catching up to the segment that can least afford \$500 million surprises.


SHARE OF WALLET, BY COMPANY SIZE

Native AI Spend: 75% Native AI Spend: 90% Average AI Spend


SMB Share of Wallet



MM Share of Wallet



ENT Share of Wallet



How does AI-Native Spend Breakdown Per Employee?

As of Q2 2026, Tropic network data puts median Native-AI ACV per FTE (full-time employee) at:




This ranking is intuitive. Smaller companies concentrate AI usage in a small number of high-usage roles, so the per-person number runs high. Enterprises spread their spend across a much larger workforce that still includes functions with no clear AI use case, and they typically negotiate better per-seat pricing as they scale.

The real story is that **Native AI spend per FTE is splitting into two different stories inside every segment:**

- 1. What the typical company spends
- 2. What the top decile spends

That split is widening, and how fast it's widening depends heavily on company size.

AI SPEND PER FTE IN Q2 2026



The growth isn't showing up in the median – it's showing up in power users, especially in enterprise.

Enterprise's median spend per FTE grew 27% year over year through Q2 2026. Its power users (P90) grew 199% over the same year, roughly seven times the median's growth rate.

Mid-Market shows the same pattern at a smaller scale: median up 10%, power users up 36%, a gap of roughly 3.5x. SMB barely shows it this year: median up 10%, power users up only 5%, though that read is skewed by an unusually high power-user quarter a year ago. Looking at the full run since Q1 2025 instead of one year-over-year snapshot, SMB's power users grew 26% against a 14% median gain, faster, but nowhere near enterprise's scale.

The gap itself backs this up. A year ago, enterprise power users spent 2.4 times the median. Today they spend 5.7 times the median, more than double the multiple in a single year. Mid-market's gap widened too, but modestly, from 2.7x to 3.4x. SMB's has held in the same 3.5x-4.4x range for six straight quarters, no clear trend either way.

The number worth watching next quarter is whether the rest of the enterprise segments starts closing the gap with its own power users, or whether that gap keeps widening.



Why enterprise, and why now?


A few things likely compound here:

Adoption in large organizations tends to happen function by function rather than company-wide, so a handful of teams with a proven use case scale hard while most of the org hasn't started.

Usage-based pricing lets those teams scale spend with consumption rather than seat count, so a function that's found real ROI can grow spend well past what flat per-seat licensing would ever allow. And the accounts showing up at P90 are likely the ones that already secured executive sponsorship and real budget, while the median enterprise account is still evaluating.

SMB doesn't show the same widening gap because it's closer to a ceiling already, most of the segment adopted early and has less room left to pull apart.

AVERAGE AI-NATIVE SPEND PER FTE, BY COMPANY SIZE



How is AI Pricing Changing?

Which vendors are losing wallet share to AI, and which aren't?


The bundle vulnerability

DocuSign's penetration within Tropic customers fell from 57% to 47% over two years, Q3 2024 through Q2 2026. Some of that shift is going to a cheaper competitor. Some of it is getting cut entirely. Both come from the same root cause.


DocuSign, Adobe, Slack, Figma, and Datadog all built their pricing on the same model: a core job plus a long list of adjacent features that justified paying a premium. Slack fell from roughly 57% to 34% penetration over the same window, the steepest decline in the group. Adobe slipped from about 51% to 45%. Figma and Datadog each moved from the high 30s and low 40s down to roughly 34%.

AI has made a lot of that adjacent stuff free. Built-in drafting, formatting, summarizing, and basic automation are table stakes now, not reasons to pay more. What's left is the narrow core job, and the narrow core job is exactly what a leaner, cheaper competitor can do well enough, or what a team can now just do with AI directly.

WHO'S GETTING SQUEEZED BY OPENAI & ANTHROPIC?



WHO'S STANDING STRONG AGAINST OPENAI AND ANTHROPIC?



AI-native crossed 34% (that's Microsoft's footprint now)

We pulled customer penetration across thousands of companies on Tropic for this window, and the pattern shows up everywhere. OpenAI and Anthropic went from essentially nothing in Q3 2024 to a combined 34% penetration by Q2 2026, the same footprint as HubSpot and Microsoft hold today. Two years ago these vendors weren't a line item. **Now they carry the same weight in the average tech stack as two of the most established platforms in the Enterprise.**

Systems of record aren't shrinking

Salesforce didn't move in that same window, holding in the mid-60s throughout. Atlassian, GitHub, and Okta stayed flat in the low-to-mid 40s. That stability isn't a coincidence, it's structural. There's no bundle to strip out of a system of record. The CRM, the identity layer, the code repository, that is the entire value. AI doesn't shrink what a team needs from Salesforce or Okta. Right now, it just logs in and uses them.

What this means for renewals

If you're renewing DocuSign, Slack, or Datadog this year, don't just ask whether AI replaced the task. Ask how much of the platform you're actually still using. A lot of teams are paying full bundle price for a fraction of the product and haven't noticed yet.

Are we still seeing an AI Tax in 2026?

Across recent renewals, we're seeing vendors run a similar [AI-tax playbook](#) we flagged in late 2025 – using AI features as the pretext for repricing. A few variations this time:



Forced bundling, no opt-out

Multiple vendors are continuing to migrate customers off legacy tiers/SKUs onto new AI-inclusive packages at renewal, with the AI features baked in rather than offered as an add-on. In several cases customers had no ability to decline the AI component even when they had no interest in using it. The price increase was as much "reduced volume" or "repackaging" as AI, but AI was the lever used to justify it.



Discounting is getting tighter

Where discounting used to be flexible (5–20%+ depending on volume), we're now seeing narrower bands (0–7%) on the usage/token components, specifically the consumption layer is where vendors are holding the line, even as they'll still negotiate on the base license.



Net uplift

Still landing in the 10–30% range, consistent with what we've tracked before — well above the standard 3–9% annual renewal increase, whether framed as a straight price hike, a "repackaging," or a reduced-volume adjustment.



Pricing model overhaul, not just price increase

The bigger shift isn't the percentage increase, but the structure. A majority of vendors are abandoning straightforward per-user pricing in favor of token, credit, or consumption-based models:

- **Usage-based/API-token pricing** replacing flat per-request or per-license pricing, with large commitment thresholds now used to waive license fees entirely for high-volume customers.
- **Credit-based AI add-on suites** layered on top of existing per-user or per-headcount pricing, rather than folded into the base price.
- **AI agent usage** bundled as an included credit allotment with overage charges once exceeded, plus new professional services/onboarding fees tied specifically to the AI transition.

One Notable Wrinkle

Early-mover discounts: Vendors sunsetting legacy tiers are offering meaningfully better pricing to customers who upgrade before forced migration deadlines. This is essentially a reward for not waiting until the renewal gun is to your head. While not a completely new negotiation lever, it's worth noting again, that customers who engage early are getting real savings, while those who wait for forced migration get list price.

How Should You Negotiate Around Credit-Based AI Pricing?

Use the forecast gap as leverage, not just a risk to manage. If this is your first year on a new usage metric, an accurate forecast is close to impossible. And because so many of these deals are moving to credit or consumption-based pricing, that forecasting gap translates directly into price unpredictability, for both sides. Vendors know this too. Say it plainly at the table. Neither of you actually knows what usage will look like in twelve months, and that uncertainty is a legitimate reason to push for better pricing and stronger commercial protections, not something to quietly absorb.

Neither party actually knows what usage will look like in twelve months, and that uncertainty is a legitimate reason to push for better pricing.

If you're moving to a credit-based model, negotiate protection on both sides of your estimate.

Underestimate usage and you're stuck paying overage rates. Overestimate and you've paid for credits you'll never use.

A few terms worth pushing for in every credit-based contract:

1 | The option to buy additional credits at a discount, instead of defaulting to overage pricing

Overages are almost always billed at list rate or worse. A pre-negotiated discount on supplemental credits protects you from penalty pricing the first time your estimate runs low.

2 | Credit rollover

Unused credits should carry into the next term instead of expiring at renewal, especially in year one, when your usage number is closer to a guess than a forecast.

3 | Overage forgiveness, up to a percentage of your total credit allotment

25% is a strong ask. Even 10%, or any buffer at all, is worth having. It's the difference between a rounding error and an unplanned conversation with your CFO.

How Do OpenAI's and Anthropic's Pricing Models Actually Differ?

The foundational AI providers deserve specific treatment because their commercial approach is unlike legacy software and unlike most other AI-native vendors.

Both OpenAI and Anthropic started with seat-based licenses. ChatGPT seats, Claude seats. Both have moved steadily toward consumption-driven pricing, and both are now in a position where they will waive seat costs in favor of driving consumption commitments.

While there's still organizational pressure not to slow down teams that AI makes more productive, that blank check is shrinking. Open-ended commitments are getting real scrutiny from finance now, not a rubber stamp.

The pricing follow-the-leader era is over

For years, ChatGPT set the market price for Enterprise AI seats and Anthropic matched it. When ChatGPT Enterprise launched at \$60 per seat in Q1 2024, Claude's Enterprise tier priced identically. When ChatGPT dropped to \$40 in Q1 2025, Claude followed again.

That pattern broke in Q3 2025. ChatGPT moved to a pooled credit model, and for the first time, Anthropic didn't match the move. Instead, Anthropic went the opposite direction in Q1 2026 with a lower seat price, around \$20, with usage priced separately on top. The license covers the seat. It does not cover consumption.

Then in June 2026, ChatGPT Enterprise adopted a strikingly similar structure. For the first time in the history of these two products, OpenAI followed Anthropic's pricing model instead of the reverse.

Both are now in
a position where
they will waive seat
costs in favor of
driving consumption
commitments.

🔗 Why this matters for buyers

The market is no longer a single price leader that everyone else shadows. Anthropic set precedent this cycle, and both providers are now converging on the same structure. They both have a lower entry price on seats paired with unbundled, uncapped consumption spend behind it. These seats are also often free when you commit to a lot of API spend. You have to be able to monitor your API usage, and benchmark rates to get ahead of it. Or, you're going to end up with big budget surprises.

Is Anthropic's model pricing actually cheaper than OpenAI's?

We tried to write this section multiple times, only for Anthropic and OpenAI to introduce new models with different costs. So, this analysis is as of August 3, 2026.

With that, there are new model generations from both providers now on the market, and new pricing. When we first started writing, Anthropic was priced slightly below OpenAI for your more basic models. But with the release of Soll, OpenAI took a big edge on Frontier models, and then they also cut prices for smaller models on July 30, 2026, as businesses continue to scrutinize AI spending more closely.

AI MODEL PRICE COMPARISON, BY TIER

Tier	OpenAI Model	Anthropic Model	Blended cost*
High-volume	GPT-5.6 Luna	Haiku 4.5	OpenAI \$0.45 vs. Anthropic \$2.00
Balanced default (through Aug 31)	GPT-5.6 Terra	Sonnet 5 (intro pricing)	OpenAI \$4.50 vs. Anthropic \$4.00
Balanced default (from Sept 1)	GPT-5.6 Terra	Sonnet 5 (standard pricing)	OpenAI \$4.50 vs. Anthropic \$6.00
Complex Production		Opus 5	Anthropic-only tier (for now) - \$10.00
Frontier	GPT-5.6 Sol	Fable 5	OpenAI \$11.25 vs. Anthropic \$20.00

*Blended cost assumes a 75% input / 25% output token mix, a reasonable default for most production workloads. Adjust for your own ratio before comparing vendors.

Two structural differences matter more than the sticker price:

Output-to-input cost ratio

OpenAI prices output tokens at 6x input across Luna, Terra, and Sol. Anthropic prices output at 5x input across Haiku, Sonnet, and Opus. Any workload that's output-heavy, long-form generation, code output, detailed agent

Token efficiency

A lower per-token rate doesn't translate to a lower total bill if the same input and output require meaningfully more tokens to represent. Run your own workload through both models before committing to a vendor based on list price alone.

The headline takeaway for buyers

Don't negotiate or budget off the published per-token rate in isolation. Price is going up for Anthropic's default tier in six weeks, the two providers charge output tokens differently relative to input, and the token count itself isn't apples-to-apples between vendors. Model your actual workload, not the rate card, before deciding where a given task should run.

Is Anthropic pushing customers into bigger commitments?

Anthropic specifically has started pushing customers toward commitment levels that mirror their current run rate or higher, and in some cases threatening to pull discounting entirely if a customer won't commit at that level. It's a page from the AWS Enterprise sales playbook, and it calls for the same pushback.

We all know your last 30 days of spending, times 12 is not a forecast. It's a snapshot of a market that's changing month to month. Hold that position directly with Anthropic. If your risk tolerance caps out at \$8-10M, say so plainly, and hold the line that Anthropic doesn't have better visibility into your roadmap than you do.

One nuance that makes this urgent rather than theoretical is that once you burn through your committed spend, even with time left on the term, you move to list rate immediately. There's no grace period. That makes it worth tracking consumption against commitment on a standing cadence, not just at renewal, so you can start an early-renewal conversation before you fall out of discount rather than after.

Where do buyers actually have leverage on AI pricing?

Foundation model providers are largely non-negotiable on headline pricing at Mid-Market volumes, but "non-negotiable" doesn't mean there's nothing to negotiate. Customers with identical commit sizes routinely land different discount rates, sometimes a percentage point or two apart, sometimes more. That gap is where the work is.

Contract terms are running short. One-year commitments are close to universal across both providers, and that's the right call for buyers too. Commit conservatively against your forecast, not to it. If you think you'll spend \$1M next year, don't commit \$1M. The model that's best for your use case today may not be the best one in six months, and a maxed-out commitment with one provider means you can't take advantage when a cheaper, faster model ships from a competitor mid-term.



A few specific terms are worth pushing on:

Get the rate table written into the contract, not linked to a website

The standard language ties your discount to a published model price list. That list can change mid-term, and your discount just applies to whatever the new price happens to be. A rate table locked into the contract holds for the full term regardless of list-price changes.

Ask for future-model discount coverage

If you negotiate a 5% discount across today's model lineup, make sure that discount extends automatically to any new models released during your term. This is a common ask and a reasonable one, but it has to be requested; it's not standard.

Understand how overage is actually priced once you exceed commitment

Anthropic's discount cuts off entirely once you cross your committed spend; everything past that point is billed at list rate. OpenAI, by contrast, continues the discount through the end of the term. That's a real difference in financial risk, and it should factor into how aggressively you commit with each provider.

Keep workloads live with more than one provider

Splitting spend across OpenAI and Anthropic is a genuine negotiating lever in a way that multi-cloud rarely is for AWS versus GCP. Re-platforming infrastructure is expensive engineering work; switching which model handles a given use case is comparatively cheap. That asymmetry gives buyers real leverage to bring to the table, not just a theoretical one.

How Is AI Pricing Creeping Into Legacy SaaS Contracts?

The same governance problem is showing up outside the foundation model providers. As AI features get layered into tools that used to be flatly seat-priced, more vendors are shifting to hybrid models such as a per-seat fee that includes a fixed number of AI credits, with overage billed on top. A \$50-per-seat CRM can turn into a \$50,000-a-month line item once a sales team starts running AI workflows against it, and procurement now has to monitor a spend category that a seat-based renewal never required.



Push for the same protective terms here that apply to the credit-based AI-native suppliers:

Credit rollover

Unused credits should carry into the next term instead of expiring, especially in year one when usage forecasts are guesses at best.

Overage forgiveness

A buffer, often 20%, that protects you from being billed in full the first time actual usage outpaces a first-year estimate.

Discounted ad hoc credit purchases

or a built-in rate table that steps down automatically once you cross a usage threshold, instead of paying penalty-priced overages for adopting the tool more heavily.

Some suppliers are also using AI as a reason to push early renewals on legacy contracts, often without saying so directly. If a supplier that's carried a flat seat-based contract for years suddenly wants to renew early "before rates change," treat that as a signal their AI functionality has gotten expensive enough on their end that they want to reprice you before you notice.

How can you manage AI Costs?

- 1 | Commit for no longer than a year on AI foundation model contracts, and commit below your forecast rather than to it.
- 2 | Maintain live workloads with more than one provider so switching is a real option, not a hypothetical one.
- 3 | Get rate tables and future-model coverage written into the contract text itself.
- 4 | Treat consumption monitoring as an ongoing operational task, not a renewal-time exercise, since the cost of missing a commitment threshold is immediate and comes with no grace period.
- 5 | Make sure users know exactly what consumes credits and train users on this information.

[Controlling AI costs](#) is primarily an internal governance problem. Matching the right model to the right task, using batch processing and prompt caching where latency allows, and setting spend guardrails at the user level so your highest-value teams keep room to work while the rest of the organization doesn't run wild.

The only way a company is going to succeed in curbing AI spend is through a company-wide discipline around how AI gets used before the bill arrives.

Fastest-Growing Vendors in the Tropic Network: Q2 2026

We ran fastest-growing vendor rankings across the Tropic network for Q2, segmented by company size. One name landed at #1 in SMB, Mid-Market, and Enterprise: Anthropic. A single vendor topping all three segments in the same quarter is unprecedented in the history of this data.

That result is no longer the most interesting finding in the dataset. The more telling signal is who's sitting just below Anthropic: Wiz, HackerOne, Anaplan, Cast AI, Sovos, Anrok, and Chainguard. None of these are AI products. All of them are growing because AI is generating more code, more cloud spend, more attack surface, and more compliance complexity than security, finance, and IT teams currently know what to do with. AI-native tools make up a real share of the top 10 in every segment, but a meaningful portion of the list is companies whose growth is a reaction to AI adoption, not a product built on it. We think this will only continue.

 SMB	 MM	 ENT
#1 ANTHROPIC	ANTHROPIC	ANTHROPIC
2  GONG	 clay	hackerone
3  nooks	WIZ ⁺	 apryse
4  Linear	 nooks	DX
5  sigma	 rocketlane	Anaplan
6  hightouch	 Chainguard	 CAST AI
7 WIZ ⁺	 Coefficient	 CURSOR
8  Juicebox	 dbt Labs	 granola
9  aircall	 OpenAI	SOVOS
10 anrok	 Notion	 cognism
Wiz and Hightouch stand out here. Both built their reputations selling into Enterprise, and both are now showing up in SMB growth rankings as security and data scalability stop being Enterprise-only concerns and become table stakes for teams of any size.		
The Mid-Market list mixes established names, Clay, OpenAI, Notion, with a faster-rising group in Nooks, Rocketlane, Chainguard, and Coefficient. It's the segment with the clearest split between companies compounding on an already-strong position and companies breaking out for the first time.		
Enterprise is the segment worth studying closest. HackerOne, Anaplan, and Sovos are companies much of the market assumed had already plateaued, and all three are back on a top-10 growth list alongside the AI-native cohort.		

What Does This Mean for Procurement and Finance Teams?

Enterprises scaling AI adoption quickly are also scaling the operational mess that comes with it, and that mess is exactly what security testing, financial planning, and compliance tooling exist to manage. Security posture management (Wiz, HackerOne, Cast AI, Chainguard) is growing because more AI-generated code and more cloud surface area means more to test and secure. Planning and compliance tools (Anaplan, Sovos, Anrok) are growing because AI spend is landing in budgets and tax jurisdictions that weren't built to track it.

The practical takeaway for procurement and finance leaders: AI adoption doesn't stop at the AI vendor line item. Budget for the second wave of spend, the security, compliance, and planning tools that AI adoption pulls in behind it, before that spend shows up unplanned in a renewal cycle you didn't see coming.



Tropic is the intelligent procurement partner for modern software buyers, helping companies save money and move faster through real-time intelligence, expert support, and agentic execution. With \$23 billion in spend under management, \$425 million in customer savings delivered, and intelligence drawn from 100,000+ completed negotiations, Tropic is where technology becomes a buyer's market. Learn more at tropicapp.io and follow us on [LinkedIn](#).