



The Procurement Leader's Budget Season Playbook

HOW TO SUPPORT YOUR CFO
WITH THE INTELLIGENCE
AI-ERA SPEND DEMANDS



Every budget season, procurement pulls together the usual stuff – contract lists, renewal calendars, spend reports. That work doesn't go away. But if that's *all* you're bringing to the table this year, you're going to have a harder time than you think.

Here's a TL;DR of what happened: software spend got more complicated right as most of our teams got leaner, not bigger. And AI is the main reason. It didn't just add a few new line items to track. It added a kind of volatility that traditional budgeting was never built to handle – consumption pricing that moves month to month, AI features getting quietly bolted onto contracts you already signed, and adoption that often starts at the top and spreads before procurement even hears about it.

So now your CFO is trying to forecast spend that doesn't behave like the old model – flat per-seat pricing with a predictable annual bump just doesn't compare to how AI spend moves. Consumption pricing and AI-driven renewal spikes are a different animal, and honestly, most finance teams don't have the visibility to see them coming on their own.

That's the opening procurement has right now: bring the right intelligence finance, IT, and other budget owners don't have the time or the expertise to gather on their own, before the annual budget and plan get finalized.

Below are five steps worth doing. None of these exercises are new to what procurement does. What's changed is how much they matter when you run them together (instead of on their own). **Together, they answer a critical question that informs budgets: what is our software and AI spend really costing us, and where's the risk hiding?** It's worth the effort because tech spend is a top two or three line item at most enterprises – and this work gives you a short, prioritized list of **what to cut, consolidate, renegotiate, or flag as risk.**

Walk in without this data and insight, and your CFO is planning around one of the biggest numbers on the P&L with a fraction of the picture.

The Procurement Leader's Budget Season Playbook

1. Map Your Tech Stack
2. Uncover Shadow Spend
3. Benchmark Costs
4. Assess the Market
5. Forecast Overages

Note: Don't treat this as five new full-time habits. Start with whatever's driving the most spend or the most uncertainty right now – usually your top 10 to 20 contracts – and go deep there first. The goal is to have real answers on the handful of things that actually move your budget.

1. Map Your Tech Stack

 **Goal** Establish a baseline inventory to surface high-priority software and AI costs

Start with the obvious stuff: a full inventory, tagged by spend amount, renewal date, and owner. While you're at it, don't throw every AI tool under one "AI spend" bucket either. This is one of the sneakier ways AI costs creep in. A vendor you've used for years doesn't necessarily cost more because you're using it more – it might cost more because a new AI tier got bundled into your renewal whether anyone asked for it or not.

We've seen that **AI spend falls into four different buckets (cost centers)**, and each behaves differently when it comes to renewal timing, pricing, and how much leverage you actually have:

- 1. **LLM licenses:** think ChatGPT Enterprise, Claude Enterprise
- 2. **LLM APIs:** think OpenAI, Anthropic, Google Gemini
- 3. **AI-native tools:** think Cursor, Glean, Clay
- 4. **Legacy SaaS with AI bolted on:** think Salesforce, Slack, HubSpot, and the growing list of tools adding an "AI tier"

Learn more: 4 AI cost centers

Split your AI spend into four cost centers: LLM licenses, embedded API costs, AI-native point solutions, and AI premiums bolted onto legacy SaaS. Lumping it all into one "AI" line item is exactly how it gets missed.

Sorting it into these buckets is also what makes the next few steps possible.

And one thing most audits skip entirely is actually talking to the people using the tools. Your contract terms and login data tell you what exists, but they don't tell you if it's working. **Sending a short pulse survey to end users** – a few numbered-scale questions on usage, satisfaction, whether it's still solving the problem it was bought for – gives you qualitative context on if this tool is really earning its keep. It's also a decent way to catch the AI tools that got adopted on excitement six months ago and are now basically collecting dust.

Once you've mapped your tech stack, flagged the most notable upcoming spend, and gathered some qualitative feedback through pulse surveys, this simple framework turns all of it into one succinct deliverable (like a spreadsheet): whether a given tool is actually meeting the needs of the business at this current time.

- **Column 1:** List of things that the business actually needs (this could be key features, key integrations, or anything that supports the needed outcome)
- **Column 2:** Defines how the current solution solves those needs (each line item is matched up)
- **Column 3:** Defines the outcome and reason for each line item in column 1 ("here's what we're getting from it")

2. Uncover Shadow Spend

 **Goal** Find under-the-radar spend so it can be flagged, consolidated, or cut

The traditional version of shadow IT hides quietly: someone expensed a tool, and it sat under the radar until IT or finance stumbled onto it. That still happens, but AI costs are like a new version of shadow IT that tends to go the other way. It gets pushed from the top, driven by real competitive pressure and a fear of falling behind, and it spreads fast, often before procurement is even in the room.

Worth remembering, too: this isn't only an AI problem. Team-level SaaS tools slip through the same way (think a scheduling app, a niche design tool). AI just made it faster and much more common.

Here's where to actually look:

- **Cross-reference your SSO/login list against your approved vendor list:** Anything showing regular login activity that isn't on your inventory list is a candidate – start there.
- **Pull recurring charges from expense reports and corporate card statements:** A lot of shadow spend never touches procurement because it's paid for on a card, not through a PO.
- **Ask IT about API keys and browser extensions:** These are the AI-specific blind spot – usage that never generates a login at all, just a quiet monthly charge on a company card or a usage-based invoice.
- **Loop in department leads directly, not just systems:** A five-minute conversation with a VP of Sales or Marketing will surface tools that may never show up in any data pull, because the team paid for it out of a discretionary budget.

Once you find something, figure out:

- Who's using it
- Why they're using it
- Whether it's actually delivering value
- Whether it makes more sense to bring it under a real contract (with real negotiating leverage) or fold its use case into a tool you already have

Any shadow spend you find can be folded back into the 3-column framework found in Step 1 to truly determine if a given tool or feature is worth your money.

3. Benchmark Costs

 **Goal** Understand where we're paying above market and how much room there is to move on price

For most of the SaaS era, pricing was pretty simple – a fixed fee per seat that scaled with headcount. That model's breaking down. Seat counts have been flat or shrinking across a lot of the market since the pandemic hiring surge faded, and vendors are making up the difference with AI-driven price hikes, forced tier migrations, and consumption or credit-based pricing that's a lot harder to pin down than a plain per-seat fee.

Understanding why pricing models have changed is important, and it's exactly why benchmarking matters more this year. The number you're comparing against might not even be the same *type* of number it was at your last renewal. Benchmarking against comparable companies, suppliers, and SKUs is still the core move, and it's worth flagging any renewal with a steep uplift before it goes unchallenged. But how you benchmark should depend on the pricing model:

- **Fixed per-seat tools benchmark cleanly against a comparable market rate**
- **Consumption or credit-based tools need usage data first – what's actually driving the credits (or tokens) – before a benchmark number tells you anything real**

Advise your CFO

To better contain AI costs, lock in contractual protections, not just the rate. This includes protections like:

- A cap on per-credit price increases at renewal
- Protection against being pushed into a pricier 'AI-inclusive' SKU without notice
- Usage alerts or spending caps

In addition to benchmarking costs, there's a layer worth adding on top of that too: not every vendor has the same amount of room to move. This is often called **price variability – the spread between what different customers pay for the same product or tier, measured by comparing the low end and high end of what real buyers pay in the market**. Some vendors show a wide spread, meaning plenty of customers are paying meaningfully less than others for the same thing (which tells you there's real room to negotiate). Others price almost uniformly across the board, with next to no gap between what any two customers pay. No amount of negotiating skill is going to move that number much.

Knowing which kind of vendor you're dealing with changes your strategy before you even get on the call. For high-variability vendors, push on price as that's where the room actually is. For low-variability ones, shift your energy to terms instead: payment timing, auto-renewal clauses, how seat additions get priced. Trying to negotiate a uniformly-priced vendor down on rate is a good way to burn a lot of time for very little movement.

By the time you've worked through exercises 1 through 3, you should have a solid handle on your top 10 to 20 contracts (priority items) – and the savings, consolidation, or cut opportunities within them that can help your CFO rightsize the budget and forecast more accurately.

4. Assess the Market

 **Goal** Identify where tools overlap and what competitive alternatives exist if we consolidate

Conduct a scan of the supplier landscape and look for overlapping tools by category – that's the familiar part of this exercise, and it's still worth doing. The goal is a shortlist of real consolidation plays.

But go a step further this time and take a fresh look at the market itself, not just your own stack against it. Markets look pretty different than they did 18, 12, even 6 months ago. New entrants, especially strong AI-native ones who are a lot cheaper, have moved into categories that used to have the same three or four familiar players. The supplier that was the obvious pick at your last renewal might not be the strongest option anymore.

Keep two separate questions in mind here instead of blending them:

1. **Internal redundancy:** Is a legacy tool's new AI feature actually duplicating a dedicated AI tool we're already paying for somewhere else? This focuses on spend you might be able to cut.
2. **Market movement:** Has a new vendor entered this category since we last looked, changing what "best option for our needs" even means right now? This focuses on whether your shortlist is still current.

And this isn't unique to AI categories, either – traditional SaaS markets consolidate too, through M&A and roll-ups, so a vendor landscape can shift for reasons that have nothing to do with AI at all.

5. Forecast Overages

 **Goal** Anticipate the pace against AI commitments (and contracts with usage tiers) to catch budget overruns

Track consumption against commitments in real time, and forecast where you're likely to see overages before they show up as a surprise.

Consumption against commitments is where AI's unpredictability shows up the most directly, and it's also the piece CFOs usually have the least visibility into on their own. This makes it one of the more valuable things procurement can bring to this cycle.

If you don't have a modern tool that can help streamline this, here's a V1 approach you can alter to your needs to anticipate and forecast overages on critical usage commitments (which you should be checking monthly):

1. **Pulling actual usage from vendor dashboards (most vendors provide these) rather than relying on billing alone**
2. **Building a run rate from 2-3 months of actual usage**
3. **Comparing projected total vs. commitment, with a concrete flag threshold (90%+ pace with months still remaining)**

And to be clear about what this is *for* – the point isn't to become the department that slows AI adoption down. It's to make sure nobody's finding out about a problem for the first time on an invoice or during budget season.

Now a forecast by itself only gets you halfway there. Before you go to finance asking for more budget, or pushing for a hard cap, pair every overage flag with a real outcome check.

What is that consumption actually producing? Not just "we're saving time," but something you can point to, like: a shorter cycle time, a better close rate, or work that used to require another hire and no longer does.

How to measure AI's ROI

When someone says an AI investment or tool "saved time," don't stop there – keep asking "so what?" until you reach a business outcome you can actually measure, like a shorter sales cycle or a faster resolution. If you can't get there, maybe the ROI isn't as advertised or needs to be revisited.

A team about to blow through its AI commitment with no answer to that question is a completely different conversation than a team that's over budget *and* has something real to show for it.

This is also where it helps to segment before you forecast, using the same four buckets from Step 1. A single "AI overage" number doesn't tell you much on its own – it could mean a team is burning money for nothing, or it could mean a team found something that actually works and is scaling it on purpose. You can't tell the difference without knowing which team, tool, or use case is behind the number.

And this kind of forward-looking check isn't only about AI credits. Multi-year contracts with built-in escalators or usage tiers deserve the same treatment – anything that can quietly step up in cost without a new conversation is worth forecasting ahead of time, AI or not.

Give Your CFO the Full Picture, Not a Fraction of It

A free savings analysis benchmarks your top contracts against real market data, so the numbers you bring to budget season hold up under scrutiny.

[Get My Free Savings Analysis](#)





Tropic is your intelligent procurement partner, purpose-built to help modern finance and procurement teams save time, cut costs, and drive impact. Whether you're a solo finance leader or part of a larger team, Tropic combines AI agents and expert services to handle the procurement work you don't have time for — from complex negotiations to renewal headaches. Powered by over \$23B (and growing) in software spend intelligence, Tropic gives you unmatched visibility, automation, and negotiation leverage across your stack. We don't just flag problems, we solve them, transforming procurement into a strategic advantage that scales with your business. Learn more at tropicapp.io and follow us on [LinkedIn](#).