SYNAPTYX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

7 things every CFO should already be doing with their AI spend

What the 2026 numbers say separates the disciplined from the exposed

SYNAPTYX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

Ask why your AI bill tripled while the price per token dropped 99%.

Cheaper tokens made it easy to greenlight every new AI workflow, but you never capped how many calls each one makes. Token prices are down close to 99% since 2022, while usage has grown fast enough to triple the total anyway.

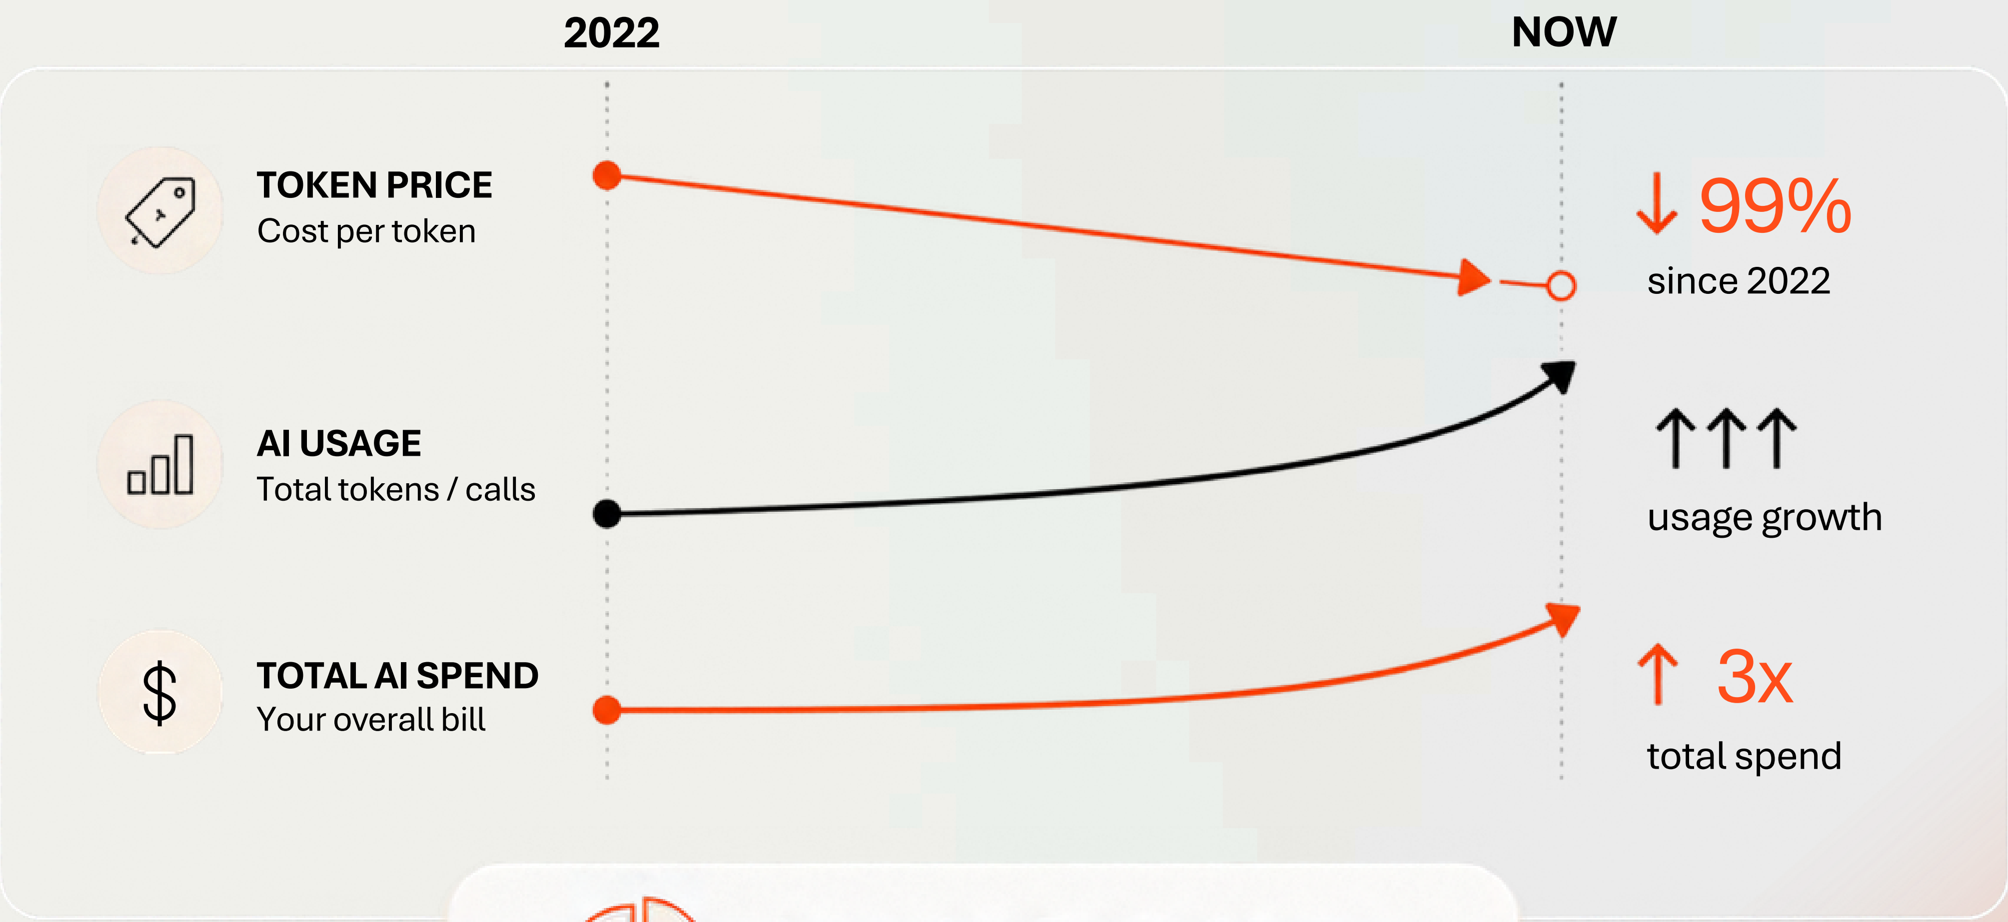

Lower prices. Exponential usage.

Higher total spend.

Someone chose your AI model, but when was the last time it was checked?

*Someone chose that model to solve one problem during the pilot, and the choice carried over as the workflow scaled to production. Costs between the cheapest and priciest models can differ by **4,500 times** per million tokens, so an unreviewed default gets expensive fast.*

CHEAPEST MODEL

Cost per million tokens

$0.02

Up to

4,500×

difference

PRICIEST MODEL

Cost per million tokens

$90.00

What made sense in the pilot can quietly become a massive cost leak in production.

Headcount clocks off while your AI agent keeps working.

A headcount line assumes fixed hours and a fixed cost. An AI agent runs whenever it's triggered, including nights and weekends, so a single busy month can blow past what your annual budget cycle planned for.

HEADCOUNT LINE

Fixed hours.
Fixed cost.

Monthly workload

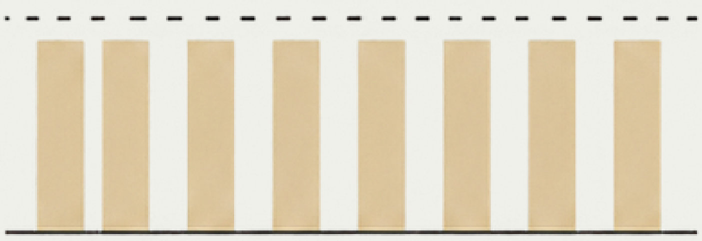

Predictable.
Stays within budget.

Always-on usage breaks the old model.

One busy month can blow past your annual budget.

AI AGENT

Runs whenever it's triggered.
24/7.

Workload over time

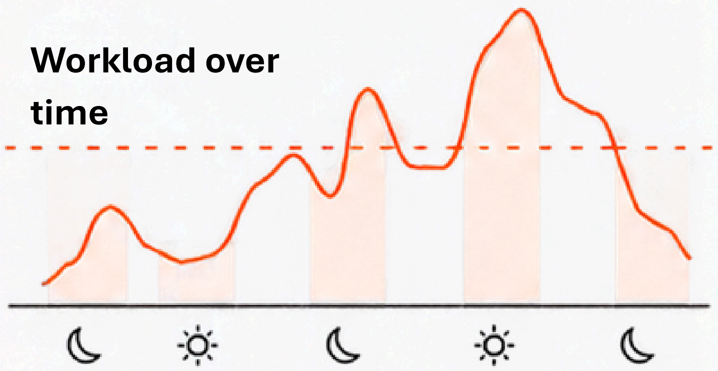

Unpredictable.
Easily exceeds budget.

AI agents don't follow office hours.
They follow demand.

SYNAPTYX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

You're funding for round two. Has round one **paid off?**

*You often approve phase two on momentum, before agreeing what phase one was supposed to prove. Bain found **44%** of large enterprises are already committing to the next round of AI spend on returns the first round hasn't delivered.*

ROUND 1: THE PLAN

- Define the problem to solve
- Run pilot & test solution
- Prove value & measure returns

44%

of large enterprises are already committing to the next round of AI spend...

...on returns the first round hasn't delivered.

ROUND 2: THE REALITY

- Momentum drives the decision
- Budget gets committed
- Value from round 1 still unproven

Momentum is not a metric. Returns are.

SYNAPTYX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

Your cloud spend gets charged back, while your AI spend doesn't.

Without chargeback, AI cost sits in a shared pool and the team driving usage never sees the bill, so they've no reason to change how they use it, even though cloud spend already gets this treatment as standard practice.

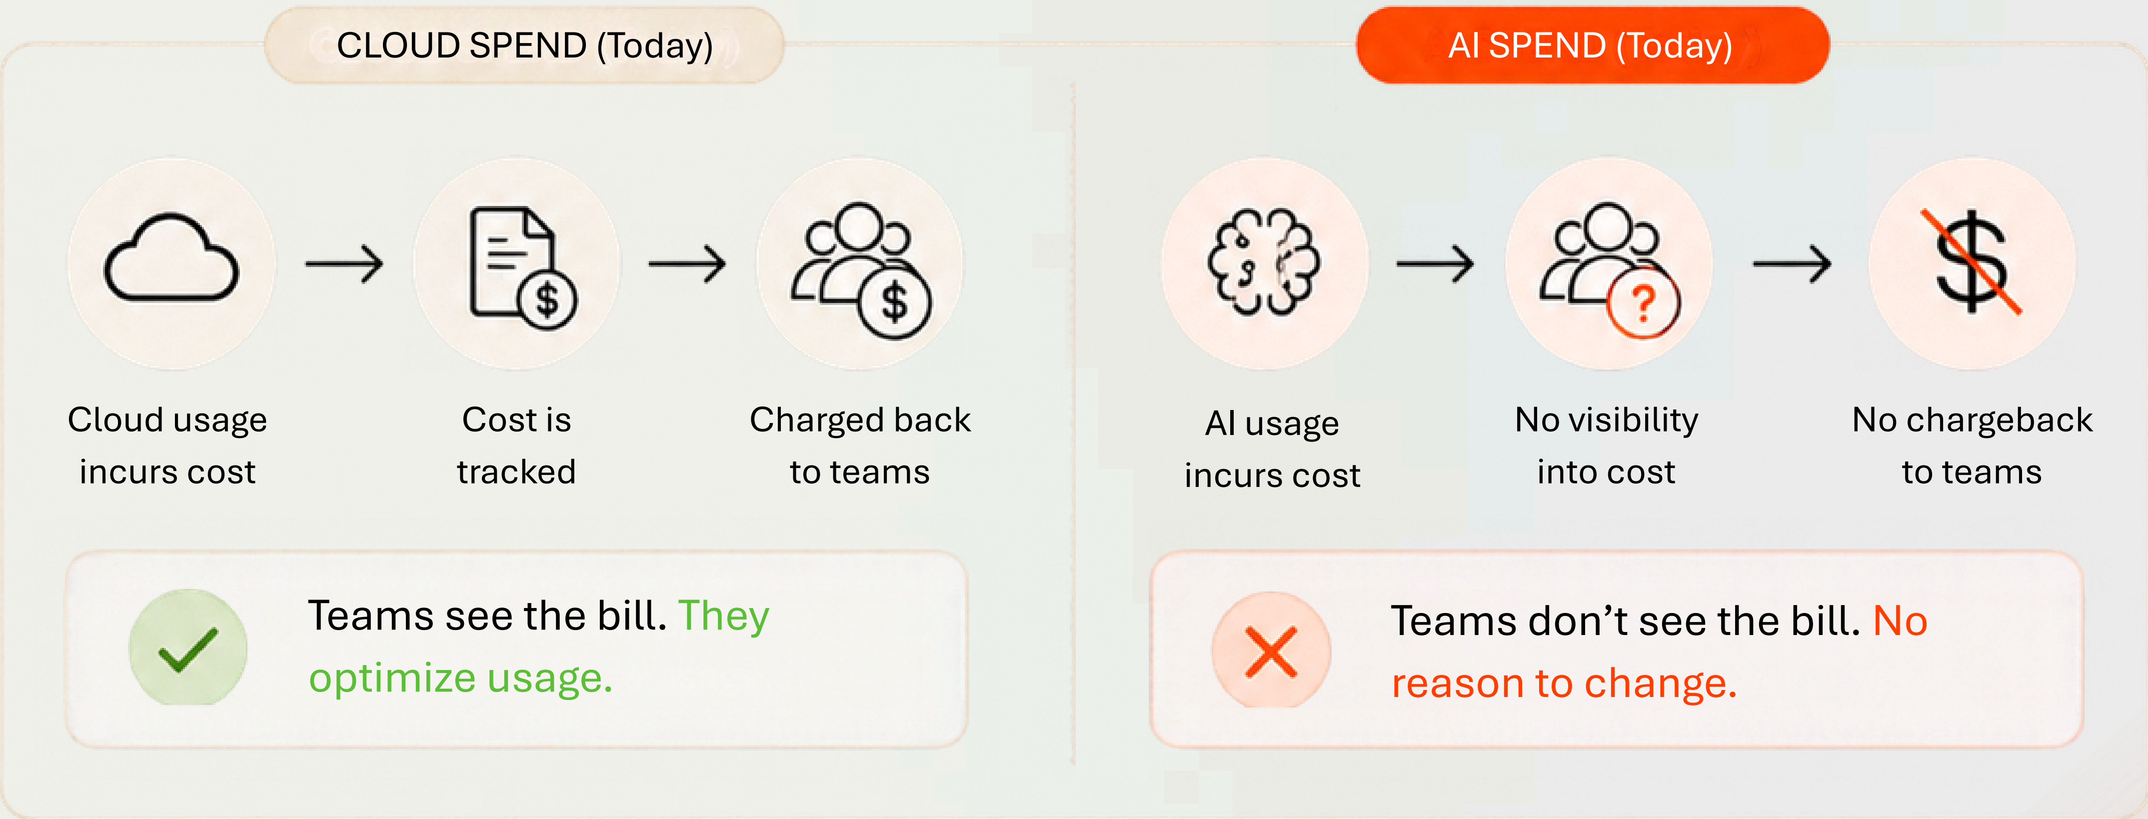

Visibility drives accountability.

Chargeback drives better decisions.

SYNAPTYX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

You've crossed your break-even point while you're still paying API prices.

Below a certain volume, API pricing wins, but cross that threshold and hosting the model yourself starts to save real money. Deloitte's modelling puts the saving at more than 50% over three years past that point, and most finance teams have never worked out where their own threshold sits.

BELOW THRESHOLD

API pricing is cheaper

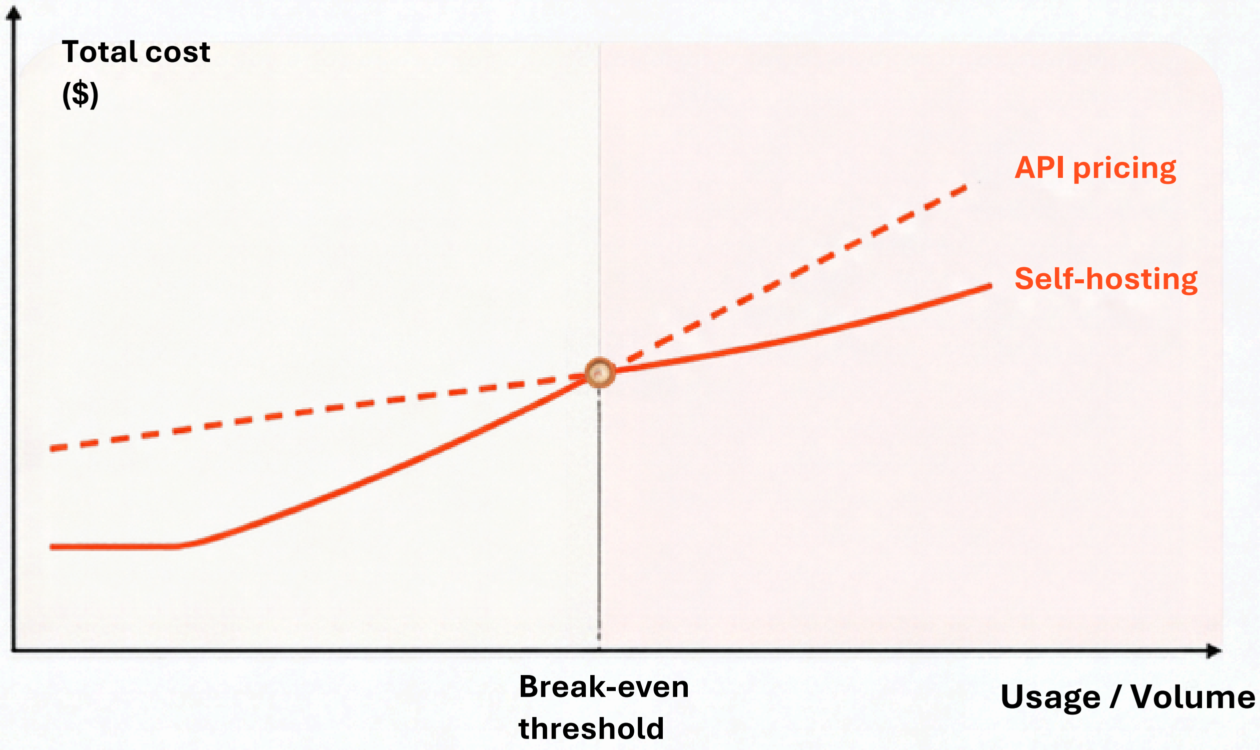

BEYOND THRESHOLD

Self-hosting saves more

Delotte's modelling shows savings of more than

50%

over three years past the threshold.

Most finance teams don't know where their threshold sits.

Know your threshold. Stop overpaying. Start saving.

SYNAPTYX

DELIVERING REAL BUSINESS IMPACT WITH GENAI

Decide whether this year's AI build is a cost or an asset, before your auditors decide it for you.

Most finance teams default to expensing AI spend as it's incurred, because no established policy exists yet for treating it as an asset. That default changes this year's margin and skews the comparison against next year, even when the work is building something the business will use for years.

TREATED AS COST

- Expensed in the year it's incurred
- Hits this year's margin
- Makes next year look worse by comparison
- No long-term value recognized

How will you treat this year's AI build?

TREATED AS ASSET

- Capitalized as a long-term business asset
- Protects this year's margin
- Reflects value in future periods
- Matches value with the years of impact

Policy today. Clarity in reporting.
Better margins. Better decisions.

Which of these seven is your business missing right now?

Tell us the number below. We read every one.

CORP@SYNAPTYX.AI