CONCIERTO | SOVEREIGN INTELLIGENCE

Sovereign Perspectives · Economics of Owned Intelligence

From Paying Consumption Bills to Creating a Strategic Asset

Rethinking AI Financials

Series Sovereign Perspectives
Category The Economics of Owned Intelligence

Format Economic Model
Published by Concierto



Version 1.0 · September 2026

Copyright © 2026 Concierto. All rights reserved.

Contents

Executive Summary

1. The Utility Illusion
2. The Cost Cliff
3. The Price That Falls, the Bill That Rises
4. Opex Without an Asset
5. Metering Work, Not Words
6. Intelligence as a Capital Asset
7. A Finance Agenda for Owned Intelligence
8. The Reframe

About This Series · Sources

In Brief

- Consumption pricing ties the cost of AI to usage, so the more value a company creates with AI, the larger its bill grows, with no ceiling it can plan against.
- Every other part of a business is paid for output; token pricing is the rare exception that pays AI for effort, and for the effort of a model the buyer does not control.
- Money spent on rented intelligence is expensed and gone each period; it builds no asset and leaves nothing on the balance sheet.
- Owned intelligence converts a volatile operating expense into a capital asset that is depreciated conventionally and, unusually, appreciates as the estate grows.
- Within a chosen tier, usage is unlimited, so a productive quarter arrives as output delivered rather than as a larger bill.

Executive Summary

The unit price of artificial intelligence is collapsing, and enterprise AI bills are rising anyway. The cost to run a model at a fixed quality fell roughly 280-fold in eighteen months, while AI moved to the center of the enterprise, with adoption jumping from 55 to 78 percent of organizations in a single year. The reason bills rise is structural: when intelligence is priced by consumption, cost scales with success, and cheaper units do not lower a bill that volume drives. This paper reframes the enterprise AI cost model, from a consumption expense that resets every quarter to a capital asset the enterprise owns and grows.

Renting intelligence by the token turns every success into a larger invoice and leaves nothing on the balance sheet.

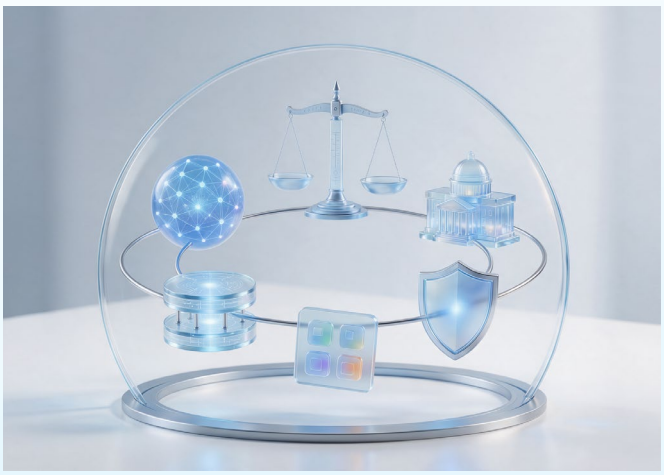

Copyright © 2026 Concierto. All rights reserved.

1. The Utility Illusion

Enterprises are buying strategic AI the way they buy electricity, metered by the unit and billed by the month. For a pilot that is sensible, because usage is small and uncertain and paying only for what is consumed is exactly right when a company does not yet know what it will use. For a capability embedded in core workflows, the same model becomes a structural mistake, and it is embedding fast: AI is no longer an experiment on the edge of the budget but a material line, and enterprise usage is climbing as quickly as the spend.

Electricity is a commodity a business consumes and never needs to own. Strategic intelligence is the capability the business runs on, and what a business runs on is worth owning.

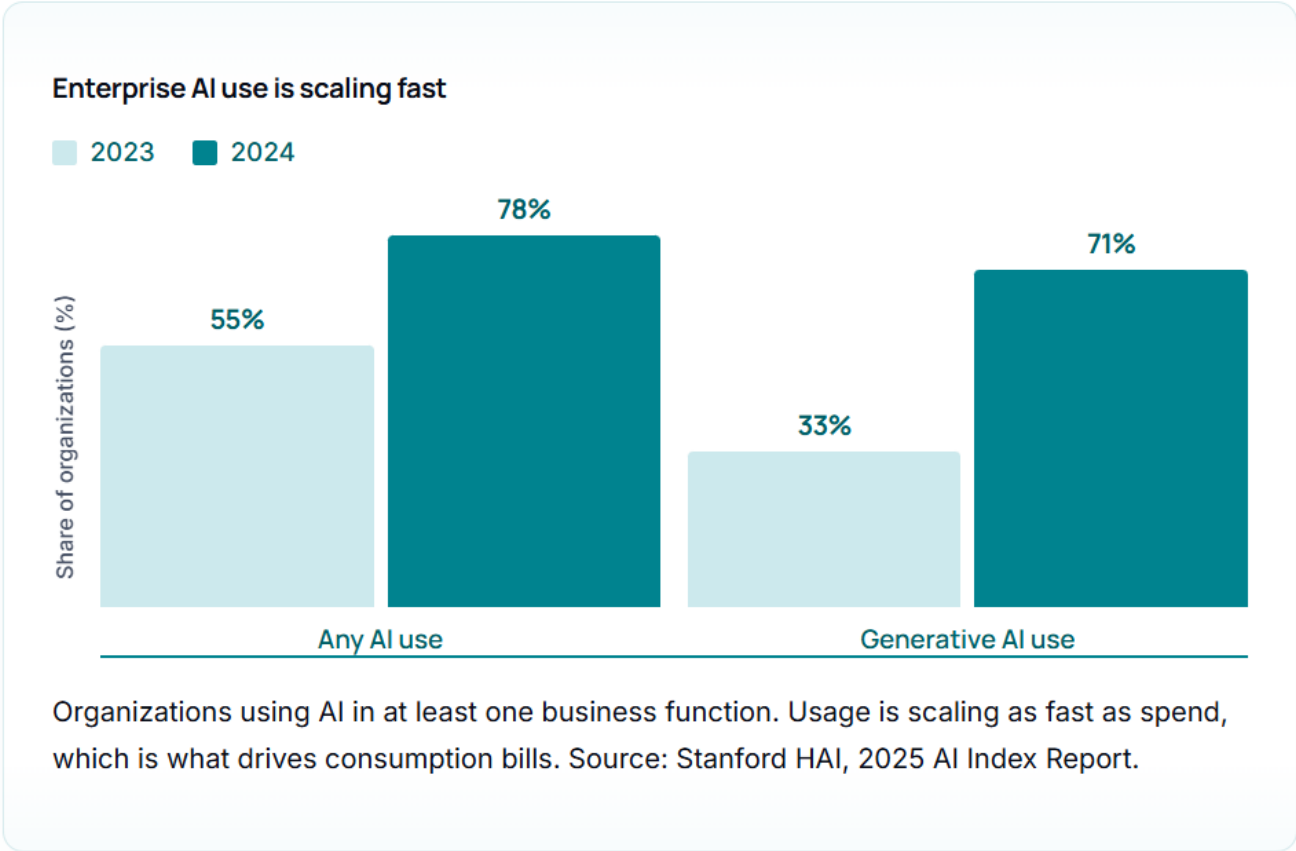

Organizations using AI in at least one business function. Usage is scaling as fast as spend, which is what drives consumption bills. Source: Stanford HAI, 2025 AI Index Report.

Figure 1. Organizations using AI, 2023 to 2024. Source: Stanford HAI, 2025 AI Index Report.

Copyright © 2026 Concierto. All rights reserved.

2. The Cost Cliff



Consumption pricing carries a property ordinary software licensing does not: cost rises with use, and it rises without a ceiling. A license fee is a known number a company pays whether it uses the software once or a thousand times. A consumption bill is a variable a company discovers at the end of the month, after the usage that produced it has already happened. The trajectory is steep and accelerating, because the newest reasoning models consume far more tokens per task than the ones they replace, and because a company adopting AI is deliberately increasing the very usage that drives the bill.

A cost that grows with success is a tax on the very adoption the enterprise is working to achieve.

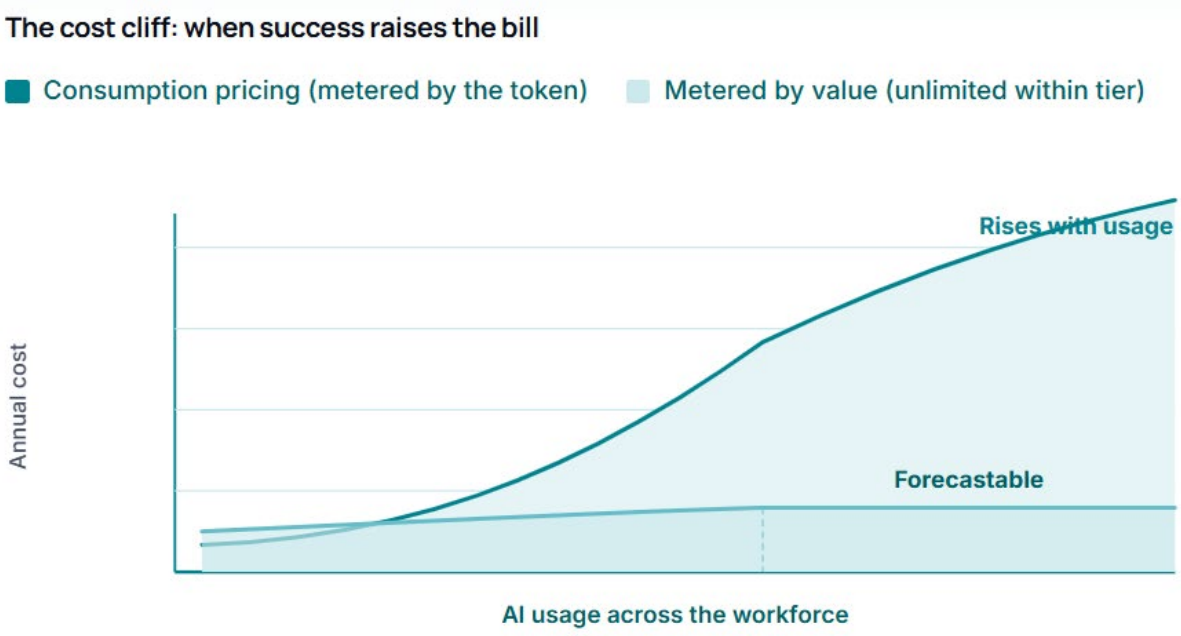

Figure 2. Consumption pricing versus metering by value. Illustrative and conceptual, not to scale.

Copyright © 2026 Concierto. All rights reserved.

3. The Price That Falls, the Bill That Rises

The obvious objection is that AI is getting cheaper, and it is, at a rate with few precedents in computing. The cost of running a model at a fixed quality has fallen by more than two orders of magnitude in under two years.

And yet the enterprise bill has moved the other way, because consumption has grown faster than price has fallen. Cheaper per token does not mean cheaper overall. A lower price for each use invites more use, and more use raises the total, so a finance leader who plans against the falling price of a token is planning against the one number in this market that does not determine the bill.

Falling unit prices are a trap for the budget, not a rescue. The number that falls is not the number that drives the bill.

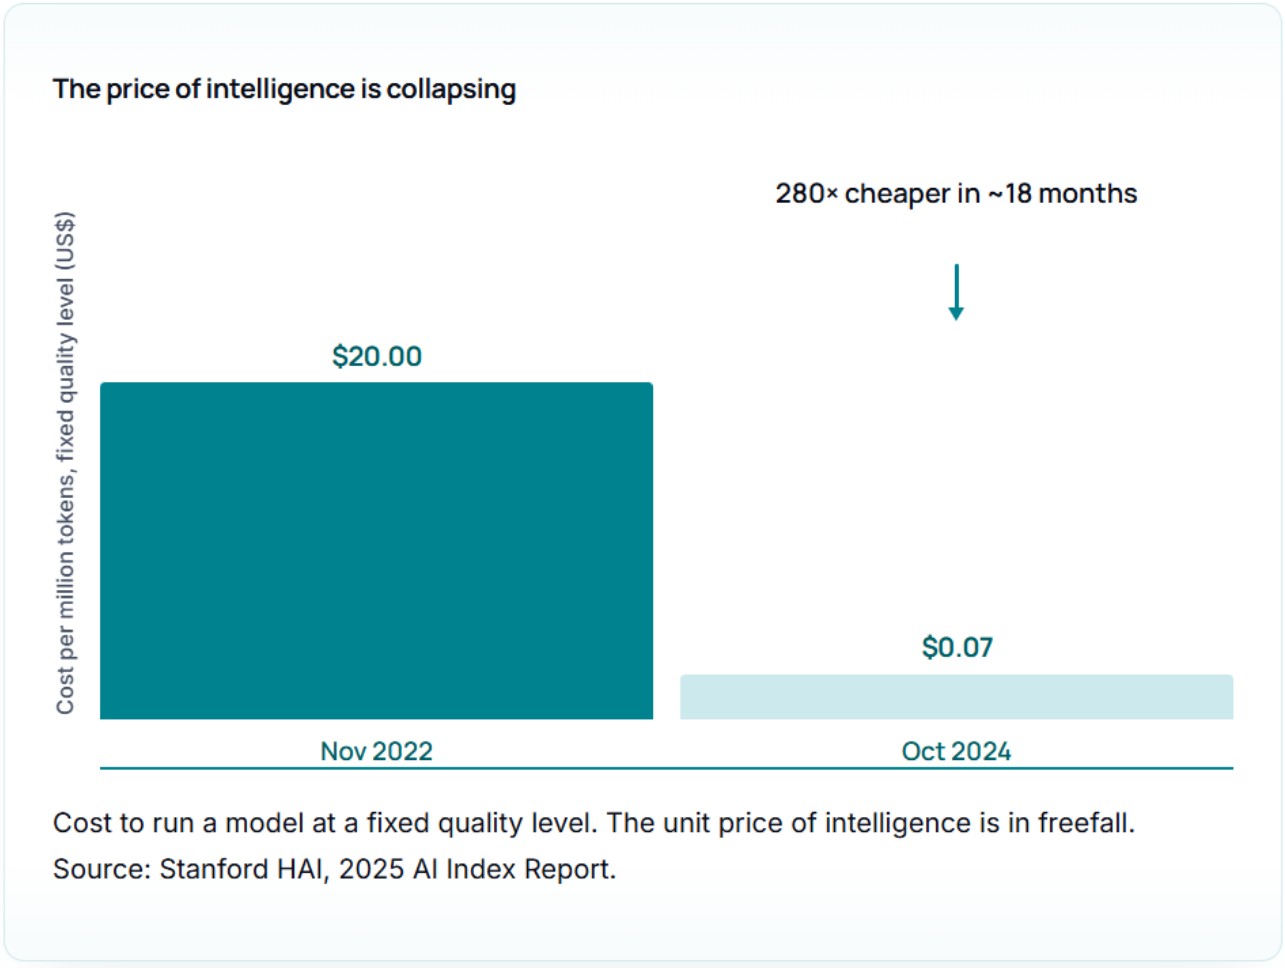

Figure 2. Consumption pricing versus metering by value. Illustrative and conceptual, not to scale.

Copyright © 2026 Concierto. All rights reserved.

4. Opex Without an Asset

Set the volatility aside, and a deeper problem remains. Money spent on rented intelligence is an operating expense. It is consumed in the period, expensed on the income statement, and gone. A company can spend heavily on AI for years, accelerate its operations, and reach the end of that spend with no asset to show for it: no model it owns, no accumulated capability on its balance sheet, only a record of tokens consumed. The investment improved a supplier's asset rather than the company's own.

This inverts a principle every other part of a business takes for granted. Companies pay their people, their suppliers, and their partners for output, for the result delivered, and not for the effort spent reaching it. Token pricing is the rare arrangement that charges for effort, and for the effort of a system the buyer does not control. The alternative is to meter the work a business actually produces, the authors, workloads, applications, and agentic events, so cost tracks the value created rather than the tokens a model happens to burn on the way there. Within a chosen tier, usage is unlimited, so a productive quarter arrives as output delivered rather than as a surprise invoice.

A company can pay in full for capability it never captured, because consumption pricing charges for the attempt regardless of the outcome.

When every part of a business is rewarded for output, there is no reason AI should be paid for effort.

5. Metering Work, Not Words

If the token is the wrong unit, the enterprise needs a better one. The token, and the vocabulary around it, serve the providers who coined them, and they obscure a fact that matters to every buyer: the model controls the quality of each answer and the number of iterations a task takes, and therefore how often the same request must be explained again and how many tokens that consumes. When a model needs three attempts to reach the answer, the customer pays for all three. The supplier earns more precisely when its model is less efficient.

Specific tier definitions and pricing are set during evaluation and confirmed by the finance office; no figures are published here.



Copyright © 2026 Concierto. All rights reserved.

6. Intelligence as a Capital Asset

The asset can be quantified. Because an owned platform carries no per-token cost, its price is close to pure margin over a compute pass-through, and it is fixed: the same seats and agents cost the same next year while the metered alternative rises with every unit of use. Averaged across a realistic estate the owned cost runs a little over half the rented equivalent in the first year, and because the owned line holds while the rented one compounds, the gap widens to roughly two-thirds by the third year. That widening gap is the return on the asset, and unlike every other capital asset it appreciates with use rather than depreciating, because more use of an owned model deepens its grounding without raising its cost.

Metering by value fixes the volatility. Ownership fixes the balance sheet, and then does something more. When an enterprise owns its intelligence, the model, the data that grounds it, and the environment it runs in, that intelligence becomes a capital asset: modeled as a capital project, carried on the books, and depreciated conventionally. But unlike an ordinary asset, it appreciates in use, because the enterprise keeps the record of that use and its unit cost falls as the estate scales rather than climbing with it. The volatile operating expense becomes a planned investment that grows more valuable the more the company relies on it.

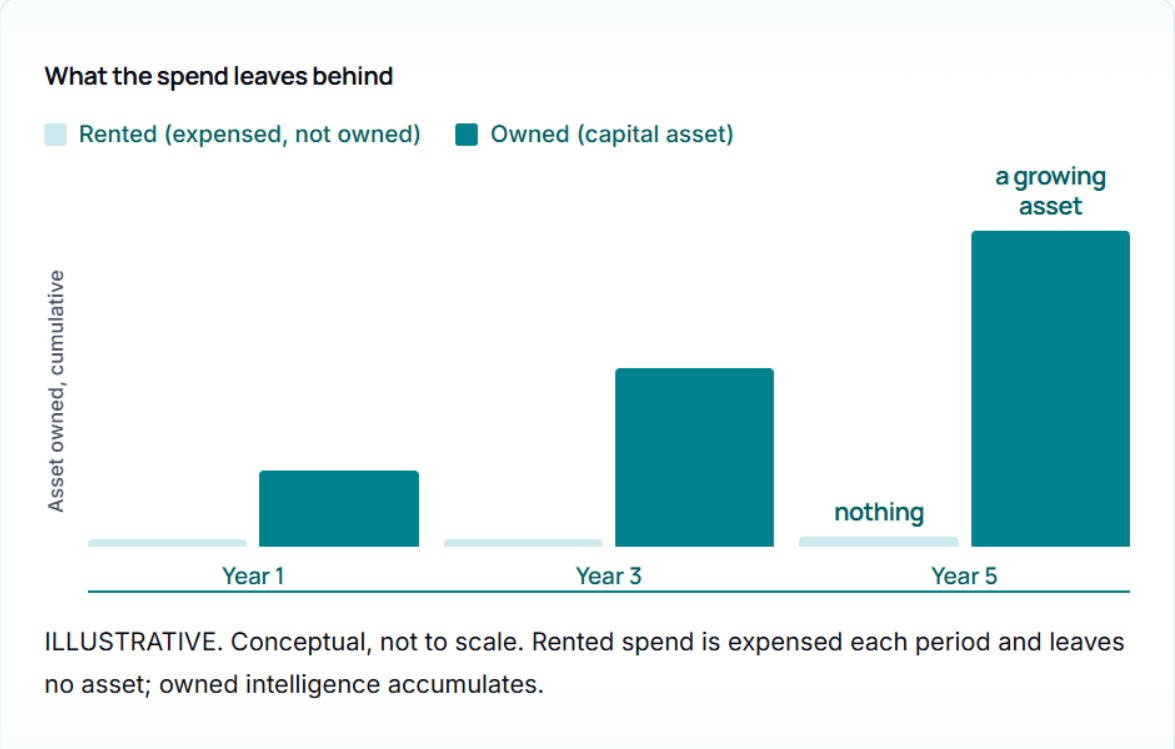

Figure 4. What comparable spend leaves behind, rented versus owned. Illustrative and conceptual, not to scale.

Copyright © 2026 Concierto. All rights reserved.

The board case follows the same discipline finance already applies to every other durable capability. The spend is already large and growing, and it is coming regardless. The only question a board actually faces is whether that money accumulates as an asset the company owns and grows, or evaporates as an expense the company rents.



Rented, AI is a cost that rises with success. Owned, it is an asset that appreciates with use. The spend can be identical; the outcome is opposite.

7. A Finance Agenda for Owned Intelligence

The reframe becomes practice through a handful of decisions a CFO and board can make now.

- **Separate the pilot from the platform.** Price experiments by consumption where that fits, and move any AI embedded in core workflows onto a model the company can forecast and own.
- **Budget on the work, not the tokens.** Adopt a unit tied to authors, workloads, applications, and agentic events, so the budget tracks value created rather than volume consumed.
- **Capitalize what is durable.** Structure owned intelligence as a capital project, carry it as an asset, and recognize that it appreciates with use rather than only wearing out.
- Test the vendor's economics, not only your own. Ask whether today's token prices are sustainable or subsidized, because a cost base built on someone else's losses will reprice.
- **Own the record of use.** Ensure the data and the learning your usage generates accrue to your own balance sheet rather than to a supplier's.

The spend is coming. The only decision is whether it leaves behind an asset or an invoice.

Copyright © 2026 Concierto. All rights reserved.

8. The Reframe

For most of computing's history the capability that mattered was owned. The cloud taught a generation of finance leaders to treat compute as a utility, metered and expensed, and for commodity compute that lesson was right. AI asks the question again, and the answer is not the same, because strategic intelligence is not a commodity a business consumes and forgets. It is the capability the business increasingly runs on, and a capability that central belongs on the balance sheet, not only on the monthly invoice.

Reframing AI from a consumption expense to a capital asset converts a volatile, unbounded operating cost into a planned investment, ensures the spend accumulates rather than evaporates, and puts the enterprise, not a meter it does not control, in charge of the trajectory of its most important cost. A company that owns its intelligence enters the next decade with an appreciating asset on its books and its budget in its own hands.

The enterprise can rent its intelligence and pay forever, or own it and be paid back. Only one of the two ends with an asset.

About This Series

Sovereign Perspectives is a research program on the ownership and governance of intelligence, published by Concierto. Its papers examine the strategic, economic, and institutional questions that artificial intelligence raises for governments and other accountable institutions, and they are written for the leaders making those decisions. The program's aim is to advance the public argument, not to promote a product.

Sources

Stanford Institute for Human-Centered AI (HAI), 2025 AI Index Report. Inference cost, corporate AI investment, and business adoption data. hai.stanford.edu.

CONCIERTO | SOVEREIGN INTELLIGENCE

reach@concierto.cloud

www.concierto.cloud

www.trianz.com

Copyright © 2026 Concierto. All rights reserved.