

FOR THE LEADER WHO HAS TO PUT ONE AI NUMBER IN FRONT OF A BOARD

Classify Every AI Consumer. Then Act on the Estate.

People and agents, across every vendor, on one graph — with nothing dropped and no owners invented.

For thirty years, every major technology spend arrived with an owner. Hardware had asset owners. Software had departments. Cloud had projects and cost centers. Generative AI broke that model, because AI is now consumed by two different kinds of consumer: employees using a dozen assistants, and autonomous agents that hold their own credentials, budgets, and decision authority.

Every serious enterprise can already see that AI is expensive. What they cannot see, in one place, is **who consumed it** — across every vendor, at the same grain, without inventing an owner where none exists.

Everyone can count tokens.
XOPS tells you whose they were.

WHAT A FORTUNE 500 ALREADY FEELS

These are not future risks. They are in this quarter's board pack.

No single number for the board

Finance has SaaS invoices, cloud lines, and card spend. None of them is "AI consumption." The board still asks for one number.

Shadow AI on personal accounts

Work leaves on consumer plans. The sanctioned estate looks smaller than reality. Risk and spend both hide.

Seats bought, seats idle

Enterprise copilots and assistants are licensed at scale. Adoption is uneven. Nobody can reclaim without looking petty.

Agent sprawl, no owner

Every team can stand up a bot. Few can name who is accountable, what it can touch, or how to turn it off.

The business will not take the chargeback

Showback fails when the dollar cannot be tied to a person, an agent, or a department without a spreadsheet fight.

Offboarding does not kill AI

Keys, seats, and agents outlive the employee. Classic joiner-mover-leaver never included this surface. It does now.

Three assistants, one job

The same persona is entitled to overlapping tools. Redundant spend looks like productivity until someone compares actual use.

An invoice is not usage

A vendor bill does not say human or agent, seat or token, or which model was actually required for the work.

The unsolved problem is the consumer, not the bill. People and agents are co-equal. Unnamed spend must stay in the total. XOPS classifies the consumer, holds what it cannot yet name, and joins the result to the enterprise graph.

COMPLETENESS

Every Dollar Stays in the Total. Even the Ones We Cannot Name Yet.

Some AI consumption resolves cleanly to a person. Some resolves to an agent. And some, on the day you look, resolves to neither — a shared key, a team account, a vendor line with no identity attached.

Every system that reports on consumption has to do something with the residue it cannot attribute. There are exactly two conventional options, and both are invisible to the buyer until the day someone challenges a chargeback.

Drop it

The totals reconcile cleanly, and they understate reality. The number presented to the board is smaller than the number the company actually spent, and nothing on the report says so.

Infer an owner

The totals hold, and the owners are fiction. A department is charged for consumption on the strength of a plausible guess. It survives exactly until the department checks.

THE THIRD OPTION

Hold it. Name it as unnamed. Keep it in the total.

XOPS holds unattributable spend in the total and labels it unnamed. It is not dropped to make the numbers reconcile, and it is not assigned to the nearest plausible owner to make the report look finished. It stays visible, and it shrinks as evidence arrives.

No invented people. No dropped dollars.

A total you can put in front of a board is one where the gaps are labeled rather than filled in.

100%

OF CONSUMPTION ACCOUNTED FOR

0

INFERRED OWNERS

WHY THIS MATTERS COMMERCIALY

Showback and chargeback only survive contact with a department head if the attribution behind them is defensible. An estate where the unnamed portion is explicit is one where every named dollar can be defended — because the system was never under pressure to invent the rest.

Named consumer, not assumed consumer. XOPS resolves a consumer to a person or an agent when the evidence is clean, and declines to resolve when it is not. The restraint is the feature.

THE CATEGORY SHIFT


One Operating Model. Two Workforces.

XOPS maintains the live AI Position of every employee and every autonomous worker: the same identity, accountable owner, declared scope, budget, delegated authority, and lifecycle primitives, applied identically to both and attributed to one org spine.

An **autonomous worker** is not a productivity tool. It authenticates into enterprise systems, consumes budget, acts within delegated authority, and triggers downstream work without a human in the loop. It can be reassigned, granted more authority, suspended, or retired. In every operational sense that matters to leadership, it resembles a worker more than an application — and most organizations still manage it as an application.

Two workforces is the reason. It is also what makes the counting possible. If there are two workforces, every unit of AI consumption ultimately belongs to one of two consumer classes — a person or an agent. Until the evidence proves which, XOPS holds it as unnamed rather than guessing. The two classes are never blended on one fact: a seat and a token are not the same object, and merging them produces a number that cannot be acted on, because no single action follows from it.

HUMAN WORKFORCE	AUTONOMOUS WORKFORCE
Employee identity	Agent identity
Manager	Accountable owner
Role	Declared scope
Budget	Budget
Approval authority	Delegated authority
Onboarding	Provisioning
Transfer	Reassignment
Offboarding	Suspension or retirement
Performance	Cost and business value



Declared outcome. Every seat and every autonomous worker stays inside its budget, inside its delegated authority, and attributed to a real owner — or is explicitly held as unnamed until it can be.

THE OPERATING MODEL

Five Things True of Every Autonomous Worker.

An autonomous workforce requires an operating model, not another technology platform. Every autonomous worker should be discoverable, accountable, authorized, bounded, and continuously evaluated. Visibility delivers the first. It does not deliver the other four.

01 · Discoverable

Found because it is consuming, not because someone registered it. A runtime consumer is not a row in a directory.

02 · Accountable

Resolved to a named human when the evidence is clean, or held as explicitly ownerless until it can be. Never guessed.

03 · Authorized

Operating inside a declared scope and an autonomy ceiling somebody consciously set, not authority nobody remembers delegating.

04 · Bounded

A budget with warn, throttle, and suspend thresholds that fire. A spending limit, in the same sense an employee has one.

05 · Continuously evaluated

Scored on identity, scope, economics, and fidelity — and on whether the work produced is worth what it costs.

Identity systems govern credentials. AI platforms manage models. Cost systems monitor consumption. Security evaluates risk.

None of them was designed to operate a workforce.

WHAT THE POSITION DECIDES

Four Questions the Position Answers. Continuously.

An AI Position is not a record of what happened. It is the live answer to who consumed what, on whose authority, against which budget, and what should happen next.

1. Who consumed this.

Every unit of AI consumption resolves to a person or to an agent, or it is held as unnamed. The two are never merged onto one fact.

2. Against which budget.

Warn, throttle, and suspend thresholds are evaluated against actual consumption as it happens. A budget is a control that fires, not a number reviewed at month end.

3. On whose authority.

Every agent and API key resolves to an accountable human and cost center when the evidence is clean — or remains explicitly ownerless until it can be resolved. Never guessed.

4. What happens next.

The Position determines the action: reassign, reclaim, throttle, retire, or fund. Because the estate is one model, the action travels across HR, identity, and finance.

ONE GRAIN

The Same Question, Asked of Every Meter.

Each AI vendor exposes its own console, its own object model, and its own idea of a user. The easy path is to model each one separately and present the results side by side. That produces a page of numbers that cannot be added together.

XOPS asks every meter the same question at the same grain, so the answers compose into one estate rather than a set of adjacent views. Parity is not coverage. Connecting to a vendor is the easy half; asking it the same question every other vendor is asked is the half that makes the totals real.

THREE SURFACES, THREE QUESTIONS

The same vendor can sell you a seat, emit activity, and bill consumption.

Those are not one number, and treating them as one is how an estate quietly stops reconciling.

Seat An entitlement someone holds. Is it used?	Token Activity actually emitted. By whom?	Dollar What was billed. Against which consumer?
--	---	---

RUNTIME IS NOT INVENTORY

A governed agent is not a row in a directory.

An application record says something was *registered*. A runtime consumer says something is *spending*, right now, under someone’s authority. Those two lists diverge the moment a team stands up a bot without registering it, and treating them as the same story is how agents go unowned.

WHAT IT SAYS	INVENTORY RECORD	RUNTIME CONSUMER
Exists	Registered by someone, once	Observed consuming, continuously
Owner	Whoever filled in the field	Resolved from the graph, or held
Spend	Not represented	Priced against a budget
Can be stopped	Only in the directory	Credential revoked, agent disabled

Coverage answers “can you connect to it.”

Parity answers “can you add it up.”

THE INSIGHT LAYER

Twelve Insights No Raw Feed Produces.

Acquisition lands the meter. XOPS decides the meaning. These are the outcomes, not how they are computed.

01 · PILLARS

Dual-pillar split

People and agents as co-equal consumers.
Never blended on one fact.

02 · COMPLETENESS

Unnamed hold

Spend with no person and no agent stays in
the total until it can be named.

03 · SURFACES

Seat, token, dollar

The same vendor can sell a seat, emit activity,
and bill consumption. Three questions.

04 · PARITY

One grain, many vendors

The same consumer question, asked of every
meter, not a new model per console.

05 · IDENTITY

Named consumer

Resolve to a person or an agent when the
evidence is clean. Never invent an owner.

06 · HYGIENE

Departed-human leakage

A key or seat that outlived the employee. A
workforce problem, not an invoice footnote.

07 · HYGIENE

Zombie entitlement

Licensed or assigned, no real use. Reclaim is
an action on the estate.

08 · AGENTS

Runtime vs inventory

A governed agent is not a directory application.
Those lists must not be the same story.

09 · AGENTS

Ownerless agent

A consumer with budget or spend and no
accountable human.

10 · HEALTH

Agent health

A scorecard on identity, scope, economics,
and fidelity. Not a spend rank.

11 · HEALTH

Role fit

What the person is entitled to versus what they
actually use. Recoverable demand.

12 · GRAPH

Cascade

Owner leaves, department showback, reclaim.
One model so one action can travel.

Twelve insights above the meter. Two pillars. Complete totals. One graph.

This is the layer a raw vendor feed cannot produce on its own, because each insight depends on a join the feed does not contain.

WHAT WE DELIBERATELY DO NOT SELL FIRST

Hop-by-hop token routing and the full agent subgraph are Full. Starter is dual-pillar identity and spend on the same estate as people, plus this insight layer. Lead with the consumer; the graph is why the consumer is worth a system of intelligence rather than a report.

NOT A CONCEPT. A SCREEN.

This Is What a Classified Estate Looks Like.

The Agent Consumption view inside XOPS Experience Center: every consumer classified, priced, scored, and attributed to an owner — or explicitly held.

my.xops.io/agent-consumption

PILLAR 02 · AGENTS · LAST 30 DAYS

AI consumption by consumer

Every meter classified to a person or an agent when the evidence is clean, priced by tokens, scored by the Agent Workforce Index, with unnamed spend held in the total.

CONSUMER SPLIT · ALL VENDORS

1,284

PEOPLE

32

AGENTS

\$14,880

UNNAMED — HELD IN TOTAL

Two pillars and the hold, on one estate. This split is the product.

ACTIVE AGENTS

32

All attributable agents cost-center owned

AGENT SPEND

\$187,305

Last 30 days

EMPLOYEE PILLAR · SEAT FIT

\$122,000

Recoverable

AGENT PILLAR · AGENT WORKFORCE INDEX

17 · 9 · 6

Green · Amber · Red

AGENTS NEEDING ATTENTION

COST CENTER

INDEX

Security Ops Agent 015

Data Analytics Cost Center

72 · Red**Support Triage Agent 010**

Customer Success Cost Center

76 · Red**Customer Success Agent 007**

IT Operations Cost Center

79 · Amber

Sorted lowest Agent Index first. Healthy agents are not surfaced here.

The split leads, not the total

People, agents, and the hold are the first three numbers on the screen. The dollar total is a consequence of the classification, not a substitute for it.

Unnamed is a state, not a gap

Held spend is displayed with the same weight as attributed spend, so nobody has to reconcile a difference between what was reported and what was billed.

Representative product view, XOPS Experience Center. Agent names, cost centers, and figures shown are illustrative.

PILLAR 01 · THE PEOPLE ESTATE

Every Employee Receives an AI Persona.

A persona is an entitlement catalog, not a budget line. It defines which models an employee may use, which tools they may reach, what their monthly budget is, and how much autonomy their AI assistant may exercise on their behalf. Limits are defined once and enforced for every employee mapped to the persona.

Executive AUTONOMY L3 MODELS Claude Opus · GPT-5.1 TOOLS Claude for Work · ChatGPT Enterprise · Microsoft 365 Copilot MONTHLY BUDGET Not set 75% WARN 90% THROTTLE — SUSPEND	Manager AUTONOMY L2 MODELS Claude Sonnet · GPT-5.1 Mini · Copilot M365 TOOLS Claude for Work · Microsoft 365 Copilot MONTHLY BUDGET \$1,000/mo 80% WARN 95% THROTTLE 100% SUSPEND
Security Analyst AUTONOMY L2 MODELS Claude Sonnet · GPT-5.1 TOOLS Claude for Work · ChatGPT Enterprise MONTHLY BUDGET Not set 70% WARN 90% THROTTLE — SUSPEND	Finance Analyst AUTONOMY L1 MODELS GPT-5.1 · Claude Sonnet TOOLS ChatGPT Enterprise · Claude for Work MONTHLY BUDGET \$700/mo 80% WARN 95% THROTTLE 100% SUSPEND

Illustrative persona configuration from a live product view. Models, tools, budgets, and thresholds are configured per organization.

Autonomy ceiling is not a model choice. It is a governance decision. Everyone talks about which models an employee may use. Very few enterprises are talking about how much authority the AI acting on that employee's behalf may exercise without a human in the loop. The autonomy ceiling makes that decision explicit, per persona, and enforces it continuously.

ROLE FIT · INSIGHT 11

Entitled, versus actually used.

The gap between what a persona grants and what its people consume is recoverable demand, quantified per persona rather than argued per person. Where the same persona is entitled to three assistants doing one job, the redundancy is visible as a comparison of use, not as an accusation.

PILLAR 02 · THE AGENT ESTATE

Meet Your New KPI: The Agent Workforce Index.

Every autonomous agent is continuously scored across four dimensions, then banded green, amber, or red. **The Agent Workforce Index is a health scorecard, not a spend rank** — an expensive agent doing exactly its job scores well, and a cheap one operating outside its declared scope does not.

Efficiency Value produced per dollar and per token. For a worker this is performance; for an autonomous worker it is cost against business value.	Identity Whether the agent resolves to a real owner and cost center.	Economic Spend against budget, fit, and expected demand.	Fidelity Whether the agent operates within its declared scope.
---	--	--	--

AUTONOMY CEILING

How much authority an agent may exercise, not just which model it runs.

L1 Drafts and recommends. A human executes every action.	L2 Executes routine actions within scope. Exceptions route to a human.	L3 Executes broadly within policy. Reviewed continuously against the Index.
--	--	---

OWNERLESS AGENT · INSIGHT 09

A consumer with budget and no accountable human.

This is a named condition with an owner queue, not an anomaly someone notices at quarter close. An agent that spends without an accountable human is surfaced the moment it is observed, routed to the nearest accountable manager in the graph, and held as unnamed in the totals until that ownership is confirmed rather than guessed.

Support Triage Agent 030	Agent workforce risk · review governance, scope, economics	Agent Index 70 · Red
Security Ops Agent 015	Agent workforce risk · review governance, scope, economics	Agent Index 72 · Red
Support Triage Agent 010	Agent workforce risk · review governance, scope, economics	Agent Index 76 · Red
Support Triage Agent 020	Agent workforce risk · review governance, scope, economics	Agent Index 80 · Red

Illustrative product view. Agent names and scores are representative of a live governance workspace.

HYGIENE

Rightsizing the AI Estate.

The same discipline that rightsizes software licenses and device fleets now applies to AI seats and agents, routing every over-provisioned entitlement to an owner for action.

Zombie entitlement
Licensed or assigned, with no real use behind it. Reclaim is an action on the estate, not a line in a report.

Role fit
What the person is entitled to, against what they actually use. The gap is recoverable demand, quantified per persona.

Redundant assistants
The same persona entitled to overlapping tools doing one job. Redundant spend looks like productivity until use is compared.

\$122,000
RECOVERABLE SPEND IDENTIFIED

200
OVER-PROVISIONED SEATS FLAGGED

100%
OF ATTRIBUTABLE SPEND MAPPED TO A COST CENTER

DEPARTED-HUMAN LEAKAGE · INSIGHT 06

Zombie AI Credentials.

API keys do not get offboarded when an employee leaves. XOPS finds every credential still active after its owner has departed, before it becomes an audit finding or a breach. This is a workforce problem, not an invoice footnote — which is why it can only be found by a system that holds the employee and the credential on the same graph.

sk-dem...9004	Microsoft	Unattributed · departed 14 days ago	Flagged
sk-dem...9003	Google	Unattributed · departed 14 days ago	Flagged
sk-dem...9005	OpenAI	Unattributed · departed 14 days ago	Flagged
sk-dem...9002	OpenAI	Unattributed · departed 14 days ago	Flagged
sk-dem...9001	Anthropic	Unattributed · departed 14 days ago	Flagged

Illustrative product view from a live governance workspace. Savings scale with estate size and persona mix.

WHY THIS CANNOT LIVE IN A REPORT

One Estate. Many Edges.

Classification is necessary. It is not sufficient. The buyer pays XOPS rates when a person, an agent, a seat, a key, a department, and a device share one model — so an action in one place cascades in another.

A person owns an agent.	An agent runs on a key.
A key outlives the person.	A seat is an entitlement , not a consumer.
A department receives showback only if the dollar can be tied to a real owner.	

Those joins are the product. Holding them on one model is what lets a single decision travel: retire an agent and the key is revoked, the budget released, the showback corrected, and the manager notified — because they were never separate records.

A spreadsheet of spend cannot do these joins.
Neither can a view built on one vendor's own object model.

CASCADE · INSIGHT 12 · ILLUSTRATIVE SCENARIO

Employee A leaves on Friday. Their agent is still spending on Monday.

Here is one edge, followed end to end. XOPS already knew Employee A owned Agent X, that it runs on API Key K, bills to Cost Center C, spends about 2.3M tokens a day against an \$8,000/month budget, and reports to Manager B.

MOMENT	WITHOUT A POSITION	WITH XOPS
Friday	Workday processes the termination. Okta deprovisions the person. The offboarding ticket closes.	The termination event fires. XOPS identifies every assistant, API key, and agent Employee A owned, with current spend and budget.
Monday	Agent X keeps running on API Key K. Credentials are valid, the \$8,000/month budget still funds it, 2.3M tokens a day keep flowing.	Manager B gets one decision — reassign ownership or retire — with inventory, spend, budget, and downstream impact shown before the choice.
Quarter close	Finance finds spend against a cost center whose owner left 90 days ago. No system tied the agent to the departed employee.	If reassigned, ownership, budget, and policy transfer across HR, identity, and finance. If retired, credentials are revoked, the agent disabled, the spend stopped.

One departure. Every assistant, key, and agent accounted for. Nothing left spending in the dark, and nothing quietly reassigned to an owner who never agreed to it.

WHAT THE BOARD WILL ASK

What Is the Maximum Financial Exposure of Any Autonomous Worker?

For a decade boards asked about cloud strategy and security posture — questions about technology investment. The next set is different. Boards are not going to debate which foundation model you standardized on; that is an implementation decision. They will ask about accountability, delegated authority, and operational control.

THE QUESTION	WHAT ANSWERS IT
How many autonomous workers operate across the enterprise?	Agent census, continuously discovered from runtime consumption
Who owns each one?	Named consumer, or explicitly ownerless with an owner queue
What decisions are they authorized to make?	Declared scope and autonomy ceiling, set per persona and per agent
What financial authority do they hold?	Budget, per worker and per cost center
What prevents uncontrolled spending?	Warn, throttle, and suspend thresholds evaluated continuously
What do they cost?	Consumption priced and attributed to the named consumer
Do they create measurable value?	Value produced per dollar and per token, scored in the Agent Workforce Index

THE ONE MOST ENTERPRISES CANNOT ANSWER TODAY

What is the maximum financial exposure of any single autonomous worker?

A human employee has an approval limit and a spending authority. Both are written down, and both are enforced. An autonomous worker consumes budget continuously and, in most enterprises today, has neither — which means the honest answer to the board's question is unbounded.

A budget with a ceiling that fires is the answer.

XOPS maintains it per worker, per persona, and per cost center.

Why this lands on the CIO. Answering it requires identity, security, finance, architecture, procurement, and business leadership to agree on one model of the estate. The CIO already sits at that intersection, which is why the autonomous workforce becomes a CIO-led operating discipline rather than a finance reporting exercise.

COVERAGE, NOT A PILOT

Enterprise Coverage, Built for the Real Stack.

The operating properties required to classify and govern AI consumption across a complex enterprise estate.

200+ENTERPRISE & AI
INTEGRATIONS**4**MAJOR AI PROVIDERS AT
ONE GRAIN**Continuous**ENFORCED, NOT ALERTED
AFTER THE FACT**Above**DEPLOYS ON THE SYSTEMS
YOU ALREADY RUN

PROVIDER	DEPTH	WHAT THE SAME QUESTION RETURNS
Anthropic	Full	Native classification, governance, and enforcement.
OpenAI	Full	Native classification, governance, and enforcement.
Google	Partial	Consumption classified; enforcement expanding.
Microsoft	File-based	File-based ingestion where no admin API exists, at the same grain.

Where a vendor exposes less, XOPS says so rather than modelling around the gap. A partial meter produces a partial answer that is labeled partial — which is the same discipline that governs the unnamed hold.

CONNECTS TO YOUR STACK

OpenAI · Anthropic · Microsoft Azure · AWS Bedrock · Google Vertex · Okta · Workday · Coupa · +200 more

CERTIFIED & COMPLIANT

ISO/IEC 27001:2022 · ISO/IEC 27018 · ISO/IEC 42001:2023 · SOC 2 Type II · GDPR

See Your AI Estate Classified for the First Time.

A 15-minute working session. No prep required.

[REQUEST A DEMO →](#)

Two consumers. One estate.
The control plane for AI Workforce Operations.

Your next workforce doesn't report to HR. It reports to XOPS.