xsparks.ai
ACCOUNTABLE FROM IDEA TO OUTCOME

AI FOR EXECUTIVES

The Software Tax

Everyone got faster. The bottom line never moved. Here is the cost hiding in plain sight.

Start with the puzzle. When ChatGPT arrived at the end of 2022, it set off the fastest technology scramble in years. Boards put AI on every agenda and told their teams to show progress. Through 2023 and 2024, mid-market companies bought what was in front of them: Microsoft Copilot seats, ChatGPT Enterprise licenses, an AI add-on for their CRM, a pilot or two per department. Enterprise spending on generative AI more than tripled in a single year, from $11.5 billion in 2024 to $37 billion in 2025 (Menlo Ventures). The tools worked: people drafted emails and summarized documents faster. And yet 56% of CEOs report no financial benefit from any of it, and only 12% saw both revenue and cost gains (PwC, 2026). **Faster people, flat margins.** Those two are not supposed to go together. They do, because of a cost almost no one put a number on. So we did.

Follow the evidence

~1,200

app switches per worker, per day. Just under 4 hours a week, about 9% of work time, lost simply reorienting.

Harvard Business Review, 2022

~60%

of the day goes to "work about work": chasing documents, status, and moving data between systems.

Asana, Anatomy of Work

100 to 275

SaaS apps the average company runs. More systems, more hours your people spend operating them.

BetterCloud / Productiv, 2024

The switching alone is roughly **200 hours per employee a year, about $8,600 per knowledge worker**, spent on nothing but clicking between systems, before you count the rest of the day spent operating them. Your finance, operations, customer-service, and sales-support teams live in this all day.

The clue most leaders miss

Here is the part that explains the flat margin. The AI did not fail; the work never left your people. Take one sales-operations analyst. A copilot now drafts her monthly report in seconds. But she still pulls the numbers out of the CRM by hand, pastes them into a spreadsheet, reconciles them against the accounting system, fixes the mismatches, and emails three managers for sign-off. The copilot sped up one step in a twenty-step process that should not need a person at all. Multiply that by every analyst, coordinator, and clerk in the building. **The tools made individuals a little faster; the work stayed in the software; the company saw nothing.** That cost has a name: **the software tax**, the hours and headcount you pay to operate software instead of running the business.

> Step back, because this is stranger than it sounds. Your people were hired to sell, to serve customers, to close the books, to make the call. Instead, most of their day goes to feeding the software that was supposed to serve them. The tools became the job.
>
> **We built a world where people go to work to operate their software, and do the work they were actually hired for in whatever time is left.**
>
> When did running the systems become the job? Once you see it on your own floor, the only question left is how fast to fix it.

What an AI-native operation does differently

This is where it changes for you. Instead of that analyst logging into the CRM, the accounting system, and three spreadsheets, she opens one place, asks for the month-end report, and AI agents pull the numbers from each system, reconcile them, draft the report, and route it for approval. She reviews it and signs off. She has gone from operating the software to directing it: the agents do the work inside the tools, she owns the result. Same person, a fraction of the time, a lower cost to run. Do that across every back-office and middle-office workflow, and the software tax comes off the books.

AI does the work. You keep the margin.

Not AI bolted onto how you run today. A business rebuilt so AI runs the work. Work on your business, not in software.

THE SOFTWARE TAX

What it costs, and what comes back

A composite estimate for a representative $250M mid-market firm, by industry.

We size the tax for the people whose day is spent in software: the back-office and middle-office roles, finance, operations, customer service, sales support, and procurement. We size it two ways, and we are clear about which is which. The **floor** is the toggling tax alone: 9% of knowledge-worker time, the figure Harvard measured. The **operating** figure is the broader share of the day spent running software, held conservatively at 35%, well under the 60% Asana found. The **recoverable value** is what an AI-native rebuild takes back, across lower operating cost and freed capacity put toward growth.

INDUSTRY	SOFTWARE-OPERATING STAFF	TOGGLING TAX (FLOOR)	OPERATING TAX (ILLUSTRATIVE)	RECOVERABLE VALUE
Professional services	~700	~$6.9M	~$27M	~$13.5M
Financial services / insurance	~550	~$5.0M	~$19M	~$9.6M
Healthcare services (multi-site)	~450	~$3.2M	~$12.6M	~$6.3M
Manufacturing	~300	~$2.3M	~$8.9M	~$4.5M
Distribution / wholesale	~280	~$2.0M	~$7.8M	~$3.9M

Composite, illustrative, and directional. The floor figure rests on Harvard's measured 9%; the operating and recoverable figures are modeled and scale with the count of software-operating staff and their loaded cost. None of it is a measured result for any one company. A tailored estimate uses your real numbers.

Worked example. A $250M professional-services firm with about 700 people whose day is spent in software. At the measured 9% switching rate alone, that is roughly $6.9M a year going to nothing but moving between systems. Count the broader operating time and the software tax runs into the mid-twenties of millions. An AI-native rebuild puts a large share of it back, on the cost line and as freed capacity for work the firm cannot staff today.

You are paying this tax today. The only questions are how much, and how fast it comes back.

Most of it is recoverable. AI does the work. You keep the margin.

Want the number for your business?

We will put it in your hands. From a few inputs, the roles whose day is spent in software, rough headcount, and how much of the day goes to operating systems, we build a tailored estimate of your software tax. It is directional, built from your numbers, and yours to keep either way. Reply and we will send it.

xsparks.ai
ACCOUNTABLE FROM IDEA TO OUTCOME

ABOUT THE FIRM

About XSparks

XSparks is a **Global AI Transformation Firm** built for mid-market enterprises that need AI to work in production. We take responsibility for outcomes, not just delivery. We fix the data, redesign the workflow, build the agents, and run them in production, composing your platforms, the wider AI ecosystem, and our own technology and methodology into one working system: the AI Operating Model (AIOM). The result is AI that delivers measured outcomes, in your operations, in your industry.

How we are different

01 **Accountable for outcomes, not just delivery.** We report the value every quarter and sign it as the operator. No AI theater, and no outcomes left anecdotal.

02 **We stay through Operate.** Most firms stop at the pilot. The return is produced after go-live, so we run the system with you, with humans in and on the loop, quarter after quarter.

03 **One operating model, not point tools.** We rebuild how the work runs so AI runs it, and instrument the value across six components: cost reduction, revenue growth, time savings, capacity gain, quality improvement, and risk reduction.

METHODOLOGY	FIRST WORKING SYSTEM	ENGAGEMENT MODEL
Think. Build. Operate.	**4 to 6 weeks**	**Outcome-accountable**

Start an Operations Briefing: a 60-minute conversation. We listen to where the operational pressure is, share what we have seen in similar operations, and decide together whether there is a fit.

CM

Cosmo Mariano
Chief Client Outcomes Officer, XSparks
cosmo@xsparks.ai · xsparks.ai

xsparks.ai
ACCOUNTABLE FROM IDEA TO OUTCOME

Think. Build. Operate. Accountable from Idea to Outcome.

Pittsburgh · New Delhi · Chandigarh · London · Dubai · Australia · Toronto

 Sources: HBR (Murty, Dadlani & Das, 2022); Asana Anatomy of Work; BetterCloud / Productiv (2024); PwC 29th Global CEO Survey (2026); S&P Global (2025).