Second-Order Effects of AI Agents Becoming Software's Primary Users

Second Order Labs™ · June 2026

ABSTRACT

When AI agents become software's primary users, seat-based pricing collapses, engagement metrics invert, and B2B marketing shifts from persuading humans to optimizing for machine comprehension.

Keywords: AI agents, SaaS pricing, enterprise software, agentic AI, business models

> *"The fundamental assumption of SaaS—that software scales infinitely while human users scale linearly—is broken when the user is an infinitely scalable AI."*
>
> — Ben Thompson, Stratechery, AI and the SaaS Mirage

Salesforce didn't launch a product at Dreamforce 2024. They hosted a wake. Marc Benioff took the stage and announced the company was "moving from copilots that require humans to agents that operate autonomously," calling it "the third wave of AI."[1] The message was blunt: the human user. The backbone of SaaS revenue for twenty years. Is being shown the door.

When AI agents become software's main users, the SaaS playbook flips upside down. Pricing models implode. Engagement metrics turn toxic. Marketing stops wooing people and starts catering to machines. The companies that see these second-order effects early will own the next era of enterprise software. The ones still counting seats will watch their business models rot.

Gartner says machine customers will drive $28 trillion in purchases by 2030.[2] That number sounds like science fiction until you follow the logic through the software stack. If the buyer isn't human, neither

is the user. And if the user isn't human, everything built for people. Interfaces, metrics, pricing. Becomes dead weight.

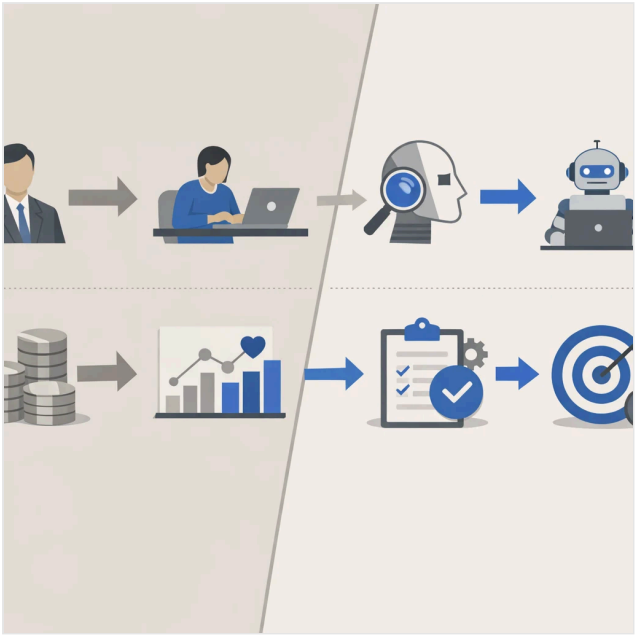

The SaaS value chain is inverting as agents replace humans at every node

I. The Collapse of Seat-Based Economics

Per-seat pricing worked because humans scale one at a time. Hire ten salespeople, buy ten Salesforce licenses. Clean, predictable, and directly tied to headcount.

Agents torch this logic. One agent with enough compute can do the work of fifty humans. Stripe's engineering team saw it firsthand: "As AI agents perform more autonomous work, SaaS companies are rapidly migrating from seat-based subscriptions to high-throughput, usage-based billing." [3]

Salesforce didn't wait to get blindsided. Agentforce launched with consumption-based pricing. Benioff called it "about outcomes, not just seats." [4] He had no choice. If one agent can replace an entire SDR team, selling seats is just selling shelfware.

Ben Thompson nailed the math: "The fundamental assumption of SaaS (that software scales infinitely while human users scale linearly) is broken when the user is an infinitely scalable AI." [5] SaaS valuation multiples depend on predictable per-user revenue. Agents blow up every financial model written since 2010.

This isn't theoretical. At Acme Corp, switching to agent-based workflows cut software costs by 30%. Tomasz Tunguz put it plainly: "Seat-based pricing is a proxy for value. When AI agents do the work, that proxy breaks down, forcing a shift to consumption or work-product pricing." [6] Companies won't buy CRM seats. They'll pay for booked meetings. Help desk software? Forget licenses. Pay per resolved ticket.

II. Engagement Inversion: When Good Metrics Go Bad

Product teams have spent years chasing Daily Active Users and session length. These numbers made sense when humans were the users. More time in the product meant more value. High DAU meant stickiness.

For agents, these metrics flip. **Engagement Inversion** is the moment when your best metrics become red flags. If an agent spends hours in your software, it's not engaged. It's stuck.

Lenny Rachitsky didn't mince words: "DAU and session length are vanity metrics for AI. If your agent is good, session length should be zero because the task was completed autonomously."[7] The ideal agent interaction? Instant. Log in, finish the job, log out. Any friction drags out the session. Long sessions mean your product is a bottleneck for agents.

The Engagement Inversion matrix shows how traditional product metrics become liabilities in an agent-first world

This flips product analytics on its head. The WebArena benchmark doesn't care about engagement duration. It tracks "task completion rate."[8] Success is about speed, not stickiness.

Harrison Chase, creator of LangChain, put it bluntly: "The UX of the future is not a dashboard; it's an agent that has the context of your entire enterprise and just gets the job done."[9] Dashboards exist for humans to make decisions. Agents don't need dashboards. They need to execute. The centralized interface becomes a relic when the user never looks at a screen.

III. The Rise of Semantic Surface Area

Agents don't watch demo videos or read case studies. They attack the API docs. Take OpenAI's GPT-4: it evaluates a product by parsing the OpenAPI spec, then runs latency tests to see if the endpoints respond fast enough. If the docs are a mess or the API is unpredictable, the agent moves on.

Semantic Surface Area is how much of your product an AI agent can actually understand and use. A product like Twilio, with clean docs and a stable API, has high semantic surface area. A legacy ERP with cryptic endpoints and spotty documentation? Invisible to agents, no matter how powerful it is.

HubSpot's marketing team saw the shift coming: "B2B buyers are increasingly using AI to evaluate vendors. If your product isn't optimized for an LLM's context window, you don't exist."[10] SEO used to mean ranking for human search. Now, agentic SEO means making your product legible to machines.

Sam Altman's line. "compute is going to be the currency of the future". Lands here.[11] Products that require fewer tokens for an agent to understand and operate will eat market share from bloated, verbose platforms. Token efficiency is the new moat.

OpenAI's function calling docs spell it out: agents interact with software through structured JSON and APIs.[12] If your product is built around a GUI, agents either have to use expensive vision-language models or skip you entirely. a16z summed up the business model: "If an AI can do the work of an SDR, you don't sell them a software seat; you sell them the completed meeting."[13]

The agentic procurement funnel replaces traditional sales cycles with technical evaluation pipelines

IV. The GUI Paradox: Why Screens Won't Disappear

It's tempting to declare the end of graphical interfaces. If agents are the users, why bother with screens? Just ship APIs and let the machines talk.

But the world isn't that tidy. Decades of enterprise software are locked behind human-centric GUIs. Rebuilding everything API-first would take years and cost a fortune. Agents need to work now, not in 2030.

Anthropic's Claude "Computer Use" feature is a direct response: "Developers can direct Claude to use computers the way people do (by looking at a screen, moving a cursor, clicking buttons, and typing text)." [14] Most enterprise software doesn't have an API for every function. So agents are learning to click through GUIs, just like humans.

The OSWorld benchmark tracks "multimodal agents capable of executing computer tasks across operating systems." [15] These agents don't care about APIs. They need screens they can read. This flips interface design: GUIs now need to be legible to vision-language models, not just pretty for humans.

Microsoft's Copilot is the poster child for this hybrid world. Their earnings call described Copilot "becoming the new UI for both the operating system and the web, acting autonomously on behalf of users." [16] The interface doesn't vanish. It morphs into a bridge between old human workflows and new agentic automation.

V. Agentic Dark Patterns and the Defense Industry

Not every vendor is rolling out the red carpet for agents. Agents optimize ruthlessly. They hunt for the cheapest path, cancel unused subscriptions, and negotiate without emotion. For vendors who profit from confusion or inertia, agents are a threat.

Agentic Dark Patterns are the new defense. These are design tricks that trip up non-human users while letting humans through.

Cloudflare saw "a massive uptick in automated AI scraping and agentic traffic" and responded by "deploying new tools to help customers manage and monetize this non-human traffic." [17] Their engineering blog details how to "block AI bots, scrapers and crawlers with a single click." [18] The infrastructure layer is now both enabling and fighting agents at the same time.

An arms race is emerging between agentic capabilities and anti-agent defense mechanisms

LangChain's documentation on multi-agent workflows doesn't sugarcoat it: "Orchestrating multiple agents requires protocols for state management, hand-offs, and human-in-the-loop authorization." [9] Human-in-the-loop isn't just about safety anymore. It's a gate vendors can use to force humans back into the loop for high-stakes actions, blocking full automation.

The law is nowhere close to catching up. Academic research on AI agent liability found that "when autonomous agents enter into contracts or cause economic harm, traditional frameworks of agency law

struggle to assign liability."[19] If an agent signs a bad contract, who pays? The company that deployed it? The API vendor? The model provider? This legal vacuum slows adoption. And opens the door for new insurance and compliance startups.

VI. The Orchestration Layer Becomes the Product

As agents take over, complexity shifts from apps to orchestration. The CRM or help desk you use matters less than the glue connecting them.

This is the cloud playbook all over again. Servers became commodities. Value moved to orchestration. Now, applications are next. The agent layer that coordinates across fragmented enterprise software is where the action is.

LangChain and Zapier are building orchestration layers that handle state across agent hand-offs, manage authentication chains, and keep context windows alive across multiple systems.[9] These aren't features of any single app. They're the new infrastructure for an agent-first enterprise.

> *"In the future, the most important customer for your business will not be a human. It will be a machine." — Don Scheibenreif, When Machines Become Customers (Gartner)*

This isn't a hypothetical. It's already happening. The real question: how fast will enterprises rip out their old assumptions and rebuild for agents?

Twenty years ago, Salesforce and a handful of others saw that software delivery could move from boxed products to hosted services. They won. Now, the delivery model is shifting again. The user is no longer human. Pricing, metrics, interfaces. Every assumption built for people is up for grabs. The second-order effects will play out over years, but the direction is set. Will your next customer even have a pulse?

Enterprise software is entering its third major paradigm shift

KEY FINDINGS

Per-seat SaaS pricing faces extinction as single agent accounts replace dozens of human licenses, forcing vendors toward consumption and outcome-based models.

Traditional engagement metrics like DAU and session length become failure signals for agents—the ideal agent session length is zero.

B2B marketing must evolve into 'Agentic SEO' where API documentation and token efficiency matter more than brand messaging.

GUIs won't disappear but must be redesigned for vision-language models as agents navigate legacy interfaces lacking API alternatives.

A liability void exists around agent-initiated contracts and errors, creating both adoption friction and new insurance market opportunities.

REFERENCES

[1] Salesforce Dreamforce 2024 Main Keynote: Agentforce

[2] Gartner, "Machine Customers Will Be Involved in $28 Trillion in Purchases by 2030"

[3] Stripe, "Usage-based billing for the AI era"

[4] Salesforce Q2 FY25 Earnings Call Transcript

[5] Ben Thompson, Stratechery, "AI and the SaaS Mirage"

[6] Tomasz Tunguz, "The Evolution of SaaS Pricing in the AI Era"

[7] Lenny Rachitsky, "How to measure product success in the age of AI"

[8] WebArena: A Realistic Web Environment for Building Autonomous Agents (arXiv)

[9] Harrison Chase, LangChain, "Multi-Agent Workflows in LangGraph"

[10] HubSpot, "LLM Optimization: How to Market to AI Agents"

[11] Sam Altman, Lex Fridman Podcast #419

[12] OpenAI API Documentation: Function Calling

[13] a16z, "The New Business Models for AI"

[14] Anthropic, "Introducing computer use, a new Claude 3.5 Sonnet, and Claude 3.5 Haiku"

[15] OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks (arXiv)

[16] Microsoft FY24 Q4 Earnings Call

[17] Cloudflare Q2 2024 Earnings Call

[18] Cloudflare Engineering Blog, "Declare your AIndependence"

[19] "Liability for AI Agents" (SSRN)

© 2026 Second Order Labs™ · https://secondorderlabs.com/articles/software-engineering/second-order-effects-of-ai-agents-becoming-softwares-primary-users/