productmind

AI INTELLIGENCE BRIEF · VOL. 01

The New Software Landscape

What the agentic shift really means for SaaS

Intelligence Brief · Vol. 01 · May 2026 · ~60 min · 22 pp. · productmind.co
· hello@productmind.co

AI is not simply a better search engine or a smarter chatbot. It is a structural shift in what software can do and who can build it. This is a field guide for software CEOs, product leaders, product thinkers & builders navigating the agentic AI era.

$37B
Enterprise AI spend
up 20× since 2023

$285B
SaaSpocalypse
wiped in one day · Feb 3, 2026

5×
AI vs internet
speed differential, Chambers

CONTENTS

What's in the Brief.

Skim the executive summary, then jump to whichever section is most relevant to the decision you are facing.

OVERVIEW

Executive Summary

We are living through what Marc Andreessen, the inventor of the Netscape browser and co-founder of a16z, recently called "the biggest technological revolution in my lifetime, bigger than the internet." This document is designed to help you make sense of that claim: what it means, what has actually changed, what parallels the internet era offers, where those parallels break down, and where the real opportunities for value creation lie.

The short version: AI is not simply a better search engine or a smarter chatbot. It is a structural shift in what software can do and who can build it.

The economics of the SaaS world that dominated the last two decades are being rewritten in real time. For a product thinker with deep domain knowledge and the ambition to build, this is one of the most significant windows of opportunity in a generation.

ON THE SPEED DIFFERENTIAL

"AI is moving at five times the speed of the internet boom, and will produce three times the outcomes of the internet age. In the internet era, a startup would spend two years building. Today, AI startups ship in weeks."

JOHN CHAMBERS, FORMER CEO, CISCO

01

Value migrates up the stack

Foundation models commoditize; vertical applications capture the durable margin.

02

The seat is dying

Outcome-based pricing is the new ground state. Incumbents structurally can't follow.

03

Domain expertise is the moat

AI capability is commodity input. Data, trust, distribution, and relationships are not.

04

Type 3 is where to play

Vertical, AI-native, built on proprietary data and earned relationships. Detailed in §6.

PART I · THE LANDSCAPE

The Landscape, What Has Actually Changed

To understand where value is created and captured in the AI era, you need to understand the new stack: the layers of technology that sit on top of each other, from raw hardware to end-user products.

1.1 · The stack, redrawn

FIGURE 01 · THE STACK, REDRAWN

The internet stack and the AI stack, layer by layer

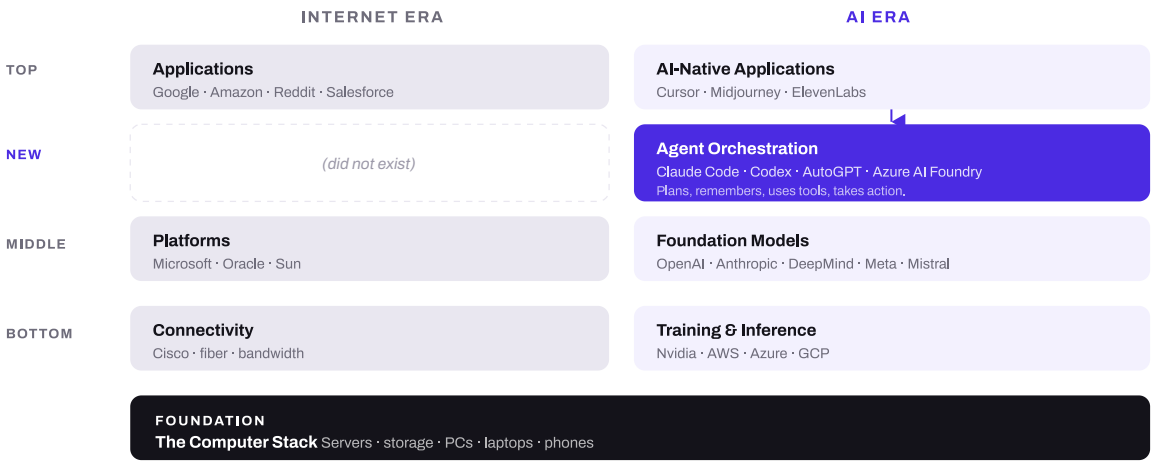

The internet era had a clean three-layer stack on top of the prior computer foundation. The AI era has a similar structure but with **a critical new middle layer**: the agent orchestration layer takes the models and makes them useful business tools that approximate human employees.

THE PATTERN

The agent layer is the new battleground. It is where software stops being a tool humans use and becomes a worker that does the job itself.

1.2 · From information to action

The defining difference between the internet-driven era we're exiting and the AI era is the shift from *connecting information* to *enabling action*.

The internet was fundamentally a medium of access and communication. It allowed people to find information faster, communicate across distances, and transact without physical presence. The human was always in the loop, reading, deciding, clicking, buying.

AI, and specifically agentic AI, removes the human from some of that loop for an expanding class of tasks. An AI agent does not just retrieve information about contract terms, it reads the contract, identifies issues, drafts a redline, and sends it for review. It does not just show you a dashboard of project risk, it emails the subcontractor, updates the schedule, and flags the budget variance.

FIGURE 02 · PACING MODEL

Era cycle times keep shortening

The technology cycle keeps shortening. Each new era arrives on top of the prior one, and resolves faster. **A two-year wait in AI today is the equivalent of a five-year wait in 2001**, long enough for the window to close.

THE CATEGORY CHANGE

This is not an incremental improvement. It is a category change, especially if you make the agent self-improving.

1.3 · The speed differential

John Chambers, former CEO of Cisco, who navigated both the internet boom and bust, puts it directly: AI is moving at five times the speed of the internet boom, and will produce three times the outcomes of the internet age. In the internet era, a startup would spend two years building a product before taking it to market. Today, AI startups ship in weeks and reach product-market fit in one to two quarters.

The advantage goes not to those who wait for certainty but to those who ship with conviction early. A startup that waits two years to validate its idea in AI may find the window has already shifted to a different shape of problem entirely.

WHAT THIS MEANS FOR BUILDERS

Speed of learning is now the strategic primitive, more decisive than capital, team size, or first-mover advantage.

PART II · THE NUMBERS

The Numbers, Capital Surge

The scale of investment into AI is unlike anything in the history of technology. The capital is real; the question is whether the revenue can sustain it.

2.1 · The capital surge

$37B

Enterprise AI spending

up from $1.7B in 2023, a roughly 20× increase in two years (Menlo Ventures, 2025)

60%

of global VC dollars

flowed to AI/ML deals in 2025, first time more than half of all VC has gone to one category (PitchBook)

$40B

OpenAI single round

the largest AI funding event in history

$13B

Anthropic, Q3 2025

29% of all AI funding that quarter

FIGURE 2.1 · CAPITAL ALLOCATION

Where the $250B in 2025 U.S. VC went

Capital is pooling at the top, foundation-model leaders command the largest checks — but the application layer is catching up.

Enterprise AI spending and global VC allocation, 2023–2025. Source: Menlo Ventures, PitchBook, ProductMind analysis. This is not distributed evenly, but **more than half of enterprise AI spend in 2025 went to AI applications, not infrastructure**.

This is not distributed evenly. Capital is pooling at the top, foundation model leaders command the largest checks. But the application layer is catching up: more than half of enterprise AI spend in 2025 went to AI applications, not infrastructure.

2.2 · Revenue is real this time

One of the key arguments against the "bubble" framing is that, unlike the internet boom, AI companies are generating genuine revenue at unprecedented speed.

10+	AI products generating over $1B in ARR globally, confirmed names include OpenAI ($10–12B), Anthropic ($5–9B by year-end), GitHub Copilot (~$2B), Cursor (hit $1B in Nov 2025, just 17 months after launch), Claude Code (~$2.5B by year-end). Salesforce Agentforce (~$1.4B in AI-related ARR) and Microsoft Copilot for M365 are plausible additions. Google Gemini's API/Workspace revenue likely qualifies but isn't broken out publicly.
50+	$100M+ ARR AI products. Per Menlo Ventures' 2025 enterprise survey, directionally credible market sizing even where names aren't published.
$2M / $4M	Median first-year ARR. Enterprise AI apps reach $2M; consumer reaches $4M (Andreessen Horowitz data).
1.5×	Faster to $5M ARR than the top SaaS firms did in 2018.
47%	AI deals reaching production. Vs. 25% for traditional SaaS (Menlo). Conversion is nearly double, demand is converting, not just exploring.
$200M	ElevenLabs revenue at the 9-month mark. Doubled in nine months to a $6.6B valuation. This is not Pets.com, users are real, demand is real, revenue is real.

THE BUBBLE ARGUMENT, WEAKENED

For all the parallels to 1999, this is the clearest disanalogy: the AI application layer is generating large, fast, real revenue today, and it is doing so while the underlying capability is still improving. Two compounding curves at once.

2.3 · Valuation premiums

AI-native companies trade at materially higher multiples than their SaaS peers.

Traditional SaaS companies now cluster in the 2.5–7× EV/Revenue range, priced as proven, stable businesses. AI startups raise at **40% higher valuations** than peers at Series A (Tomasz Tunguz, Theory Ventures). The premium for being AI-native, at least in early-stage pricing, is now structural rather than narrative.

But the market is getting more discriminating. Investors are paying for demonstrated impact on revenue, margin, or unit economics, not for AI as a story. The "AI inside" trade is closing; the "AI does the work" trade is opening. Founders who can show a step-change in unit economics, cost per ticket, cost per invoice, cost per lead, are getting funded at multiples that look more like venture pricing than public-market pricing.

FIGURE 2.2 · VALUATION

EV / revenue multiple ranges

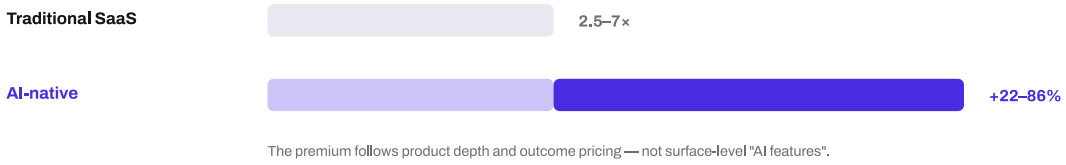

Case · Cursor · $29.3B. Hit $1B ARR in 17 months. Priced at $29.3B in November 2025, venture pricing applied to public-scale revenue. High multiples follow product depth + outcome pricing. They do not follow surface-level "AI features". (Cross-refs: §2.2 · Revenue real; §6.1 · Three types; §7.1 · Indicators.)

PART III · THE SAAS EARTHQUAKE

The SaaS Earthquake and What Is Breaking

3.1 · The death of the seat

The SaaS business model was built on a simple, elegant idea: charge per user, per month, forever. One human, one license. Revenue scaled with headcount. The entire go-to-market motion of the last 25 years was built on this logic.

Agentic AI is dismantling this logic. A single AI agent can now perform the tasks that previously required dozens of people. Why pay for 100 seats in a marketing department when three autonomous agents can execute the same volume of work through a single API?

> FEB 3, 2026 · "SAASPOCALYPSE"
>
> Approximately $285 billion in market value was wiped from global SaaS and enterprise software companies in 24 hours. Salesforce and Adobe each fell more than 25% as investors repriced the per-seat model out of existence.

Salesforce is attempting to pivot with its "Agentforce" platform, charging $0.10 per autonomous "action" via Flex Credits. Adobe faces an existential challenge as its creative moat, built over decades around specialized tools, is being bridged by generative AI from all sides.

INEXORABLE, IF SLOW

It might take a while, technology diffusion takes longer than tech insiders imagine, but this value logic is inexorable and will put significant pressure on the SaaS model over the next 24 to 36 months.

3.2 · The new pricing battleground

The industry is in an uncomfortable transition. The old model is broken; the new model is not yet standard. Three pricing paradigms are competing.

DYING

Seat-based

Pay per human user. Incumbents like Salesforce, HubSpot, and Asana are trapped here.

CURRENT

Usage-based

Pay for what you consume, tokens, API calls, tasks. The dominant model for foundation model APIs.

EMERGING

Outcome-based

Pay for results, tickets resolved, invoices processed, leads qualified. Zendesk, Intercom, Agentforce are early experiments.

FIGURE 03 · PRICING SHIFT

Seat → usage → outcome: value repositions to the work

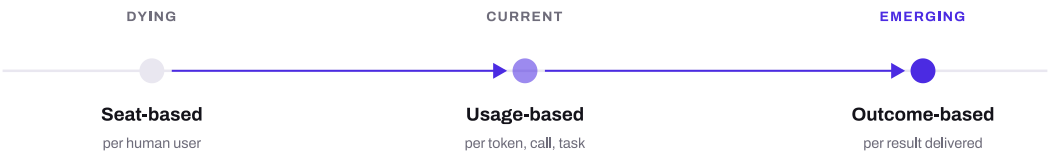

By 2028, Gartner expects 70% of vendors to refactor pricing; IDC sees 40% of SaaS spend shifting by 2030.

Gartner predicts that by 2028, pure seat-based pricing will be obsolete, with 70% of software vendors refactoring their pricing strategies around new value metrics. IDC projects that at least 40% of enterprise SaaS spend will shift toward usage-, agent-, or outcome-based pricing by 2030.

THE ASYMMETRY FOR BUILDERS

You can architect a business from the ground up with outcome-based pricing, something incumbents are structurally incapable of doing without cannibalizing their existing revenue. That asymmetry is your opening.

3.3 · The commoditization of software code

Much of what took an 18-month development cycle and a six-figure engineering team to build in 2020 can now be prototyped in days.

The cost to train a high-performance AI model has collapsed, from roughly $100 million for OpenAI's flagship model to the $5 million claimed by DeepSeek, to just $30 for Berkeley's TinyZero reproduction in 2025. Whatever you believe about the precise figures, the direction is unambiguous: intelligence is rapidly approaching zero marginal cost.

The practical implication: software itself is becoming a commodity. The moat is no longer the code. It is the data behind the code, the trust it takes to deploy it in a specific context, and the distribution to reach the right customer.

> "A lot of SaaS is essentially a form with logic on top of a database. That is now a commodity. The companies that survive own something AI cannot easily replicate: proprietary data, trust, or distribution."
>
> WIDELY ATTRIBUTED, SEQUOIA / BAIN ANALYSIS, 2025

TRAINING COST · DESCENDING

~$100M · GPT-class (2023) → **~$5M** · DeepSeek (claimed, 2025) → **~$30** · TinyZero reproduction. The numbers are contested. The trajectory is not.

3.4 · The rebundling

SaaS unbundled enterprise software into specialized point solutions, a separate tool for CRM, HR, finance, analytics, and project management. The result was software sprawl: the average enterprise uses hundreds of SaaS tools, each with its own login, data silo, and integration headache.

Agentic AI is reversing this; a rebundling is underway. Bain's 2025 Technology Report describes the emerging enterprise stack as three layers: **systems of record** (the source-of-truth data layer), **agent operating systems** (the orchestration layer), and **outcome interfaces** (what the human sees and controls). A single agentic system can orchestrate workflows across multiple tools without the human ever touching a dashboard. IDC envisions a future where users instruct

an AI agent to "approve last week's expense reports" or "generate next quarter's sales forecast", and the agent orchestrates workflows across HR, finance, and CRM systems behind the scenes, with no direct interaction with any individual SaaS product.

> WHAT THIS MEANS
>
> SaaS unbundled enterprise software. Agents are quietly putting it back together, under one outcome interface, with the human at the edge rather than the centre.

PART IV · THE INTERNET PARALLEL

The Internet Parallel, What We Can Learn From It and What We Cannot

4.1 · What the internet era teaches us

The internet boom of the late 1990s rhymes with the AI boom in important ways. Understanding the parallel helps you avoid repeating history's mistakes while recognizing the genuine opportunity.

01 · TRAP

The Infrastructure Trap

Capital pools at the bottom of the stack and most of it is destroyed. The remainder seeds the next era.

02 · WAR

The Platform Wars

Subsidized access → eventual monetization. The model layer is fighting for default position.

03 · LAYER

Application Layer Wins

The greatest value was always at the application layer. The AI era should be no exception.

FIGURE 04 · THE TAKEAWAY

Build trains, not rails: value migrates up the stack

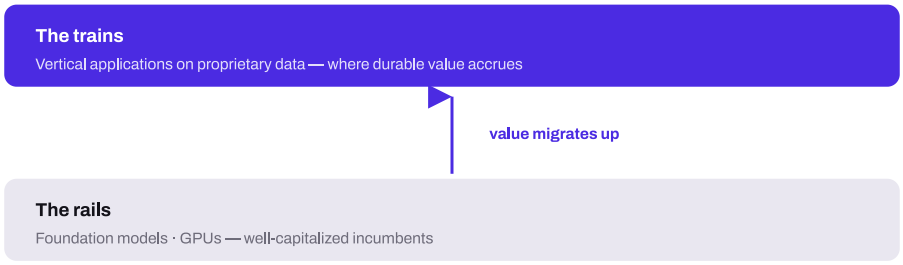

Build trains, not rails. The rails (foundation models, GPUs) are well-capitalized incumbents. **The trains (vertical applications on proprietary data) are where durable value accrues.**

4.1.1 · The infrastructure trap

During the internet boom, enormous capital flowed into infrastructure, laying fiber, building data centers, and manufacturing networking equipment. Most of those companies failed or saw their stock price fall 80–90% from their peak. Cisco, which briefly became the most valuable company in the world at $550 billion in March 2000, is still trading well below that peak today.

The same dynamic is emerging in AI. Goldman Sachs warned in 2024 that companies worldwide are poised to spend $1 trillion on AI infrastructure with "little to show for it." Sequoia partner David Cahn identified a "$500 billion revenue hole", the gap between what AI infrastructure spending implies in revenue expectations and what is actually being generated.

This does not mean the infrastructure is worthless, Global Crossing's fiber became a gift to the cloud era, it means the value migrates up the stack, to the companies that use that infrastructure to build things people actually pay for.

4.1.2 · The platform wars

The internet gave us platform wars, Microsoft vs. Netscape, Google vs. Yahoo, Facebook vs. MySpace. In each case, the winner was the company that controlled the most valuable layer of the stack, and worked to shut other companies out (the internet spawned a lot of anti-competitive activity, and anti-trust regulation followed).

The AI era has its own early platform war at the model layer: OpenAI, Anthropic, Google DeepMind, and Meta are competing to become the default intelligence substrate. The dynamics are similar, massive subsidized access to attract developers, followed by eventual monetization. But the stakes are higher and the switching costs are potentially greater, because switching your underlying model is not quite like switching your search engine.

> $500B REVENUE HOLE
>
> The Sequoia bear case. Infrastructure spend implies a level of revenue that the AI application layer has not yet generated. The hole may close as applications scale, or it may stay open and force a correction. Either way it is the single most-watched number in venture in 2026.

4.1.3 · The application layer opportunity

In the internet era, the greatest value was created not by those who built the infrastructure but by those who built the applications: Google, Amazon, Facebook, Salesforce, and hundreds of vertical SaaS companies. The infrastructure companies were important, but the application builders captured more sustained value.

The same is likely true in the AI era. The foundation models are impressive and important. But the $37 billion in enterprise AI spending is going increasingly to applications, to the products built on top of the models, not to the models themselves.

WHERE THE MONEY IS MOVING

More than half of enterprise AI spend in 2025 went to applications, not infrastructure. The center of gravity has already started its migration up the stack.

4.2 · What it doesn't teach us (where things are different)

The parallels are useful, and they have limits. The next section covers the five places the AI era is materially different from the internet era. Pattern-matching from the 1999 playbook past those five places will mislead you.

Together these are the places where "the AI era is just the internet again, faster" stops being a useful analogy and starts being a misleading one.

1. **Speed.** The internet era unfolded over a decade. The AI era is unfolding in years. A startup that waits two years to validate its idea in AI may find the window has already shifted, to a different problem shape entirely.
2. **Barriers to entry.** The internet required physical infrastructure, servers, bandwidth, datacenters, that created real barriers to entry. AI infrastructure is virtualized and being commoditized at speed. This means it is easier to build, and that your competitive moat needs to come from somewhere deeper than in prior paradigm shifts.
3. **The human displacement question.** The internet created new categories of jobs while disrupting old ones, it automated some work while expanding the workforce that could do new things. The AI era poses a more direct challenge: agents do not just assist workers; they replace workflows. The economic, political, and ethical implications differ in scale and speed.
4. **Data as moat.** In the internet era, network effects and data were emerging concepts. In the AI era, proprietary data is the primary moat. A model trained on generic internet data is a commodity; a model fine-tuned on 20 years of your industry's specific projects, contracts, and outcomes is a defensible asset.

5. **The regulatory environment.** The internet grew up largely unregulated, especially in the U.S. The AI era is arriving in a world actively grappling with how to regulate it. The EU is moving quickly with the AI Act (adopted May 2024), perhaps far too quickly given it does not largely originate this technology. The U.S. has retreated from thoughtful anti-competitive regulation and is torn between federal and state efforts. California is emerging as the most mature regulatory district in the nation.

> REGULATION AS A FEATURE
>
> Compliance costs are inevitable. But builders who treat regulation as a product feature, auditability, data lineage, jurisdiction-aware policy, turn a headwind into a moat.

PART V · THE DEBATE

What Silicon Valley Is Actually Debating

5.1 · The optimists

Marc Andreessen (a16z).

> "This is the biggest technological revolution in my lifetime. In terms of magnitude, it is obviously bigger than the Internet. Comparable things are the microprocessor, the steam engine, and electricity."
>
> MARC ANDREESSEN, A16Z

Andreessen argues that AI will be fundamentally deflationary, intelligence costs will "drop like a stone" as competition intensifies, similar to how the cost of computing, storage, and bandwidth all fell to near-zero over the internet era.

Martin Casado (a16z). While acknowledging that "private valuations are crazy," Casado professes no concern about an AI bubble. "This stuff is magic. The users are real. The demand is real. The GPU usage is real." a16z has committed a $3 billion fund specifically to AI infrastructure bets.

Ben Horowitz (a16z). On the firm's AI portfolio: "It is one of the best funds, like, I have ever seen.", while also cautioning it is too early to judge.

THE OPTIMIST'S CORE CLAIM

Intelligence is becoming as cheap as electricity. The companies that figure out how to wire it into specific real-world workflows will compound past the model labs themselves.

5.2 · The skeptics

David Cahn (Sequoia). Published the most prominent bear case in the VC world, identifying a growing revenue gap, expectations of $600 billion in AI revenue against infrastructure buildout that implies even greater spend.

> "Speculative frenzies are part of technology, and so they are not something to be afraid of. Those who remain level-headed through this moment have the chance to build extremely important companies."
>
> DAVID CAHN, SEQUOIA

Gary Marcus (AI researcher). Calls OpenAI "the WeWork of AI" and is so confident current approaches cannot achieve AGI that he bet Anthropic CEO Dario Amodei $100,000 that AGI would not be achieved by the end of 2027. His argument: without AGI, the valuations of leading AI startups stop making sense.

Goldman Sachs. Companies worldwide are poised to spend $1 trillion on AI infrastructure with limited visible return. Investment risk stems from "a handful of tech stocks accounting for an uncommonly high share of market capitalization."

5.3 · The prevailing concerns, what they agree on

- **Revenue concentration.** Too much capital going to too few players at the model layer, creating systemic risk.
- **The DeepSeek shock.** A Chinese model matching frontier U.S. performance at allegedly 95% lower cost destroyed the assumption of an insurmountable U.S. moat. The directional implication, that model intelligence will commoditize, is widely accepted.
- **ROI anxiety.** Enterprises ran pilots in 2025 and are demanding evidence that AI spending translates to outcomes. Some are seeing sticker shock: the cost of effective intelligence often exceeds the cost of human staff today.
- **Safety vs. speed.** Is the industry moving too fast for safety guardrails to keep up? Largely a regulatory/policy battle, not a product one.
- **Job displacement.** What happens when AI agents automate not just repetitive manual work but knowledge work, the cognitive tasks that constituted white-collar employment for decades?

PART VI · THE PM LENS

Where Value Is Created, The PM Lens

6.1 · The three types of AI businesses

Not all AI businesses are equal. The ones that create durable value share specific characteristics. It helps to think in three categories.

FIGURE 05 · THE THREE TYPES

From AI at the edge to specialized AI at the core

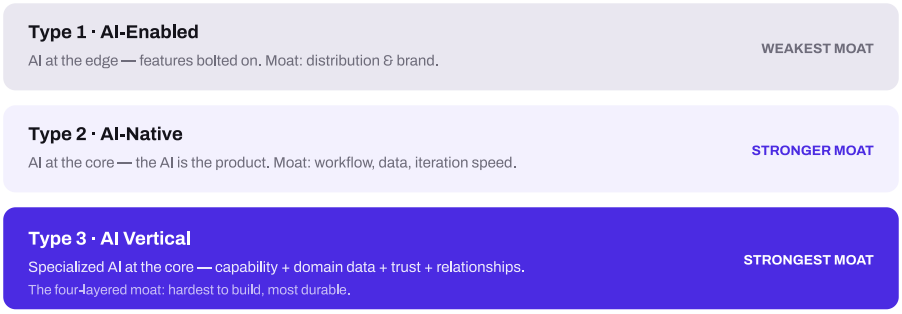

For a product leader with 10+ years of domain experience, **Type 3 is the highest-value opportunity**. The intelligence is becoming a commodity. The domain knowledge, the trust, and the relationships you spent a decade earning never will.

Type 1 · AI-Enabled, AI at the edge (weakest moat). Traditional software companies that add AI features, a chatbot on top of an existing product, AI-powered search, GPT integration in a dashboard. These are table stakes by 2026, not differentiators. The moat here is distribution and brand, not the AI itself.

Type 2 · AI-Native, AI at the core (stronger moat). Products designed from the ground up around AI capabilities, where the AI is the product, not a feature. Cursor (AI code editor, valued at $29.3B in November 2025), Midjourney, ElevenLabs. These companies can price for outcomes because the AI is doing the primary work. The moat is the workflow design, the proprietary data they accumulate, and the speed of iteration.

Type 3 · AI Vertical, specialized AI at the core (strongest moat, hardest to build). Companies that combine AI-native architecture with deep domain knowledge, proprietary data, AND earned relationships inside a specific industry. This is where the most durable value creation lies for teams with domain expertise. The moat is four-layered: the AI capability, the domain data, the trust of a specific customer base, and the relationships and credibility it takes years to build in that field. Relationships matter as much as the data, often more, because they are what get the agent deployed inside the workflow in the first place.

6.2 · The moat hierarchy in the AI era

In the SaaS era, moats came from network effects, switching costs, and scale. In the AI era, the moat hierarchy looks like this, from weakest to strongest:

FIGURE 06 · MOAT HIERARCHY

Four moats, weakest to strongest

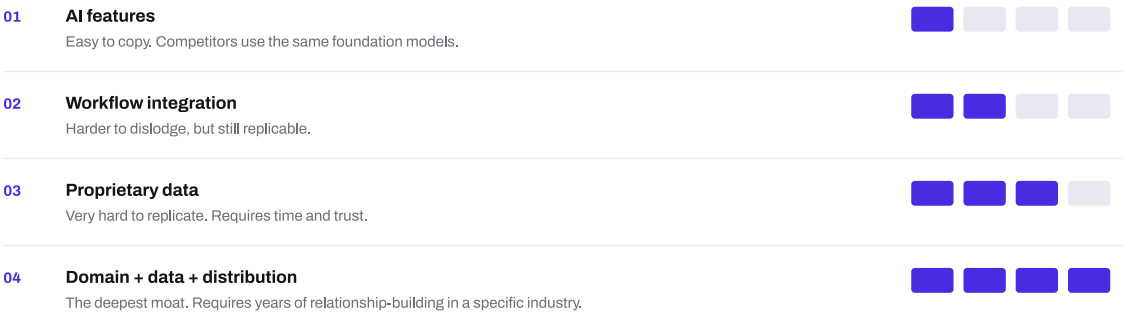

Moats in the AI era, ranked weakest to strongest. The deepest — **domain + data + distribution** — is the one an AI-native challenger cannot buy, because it is built from years of relationship in a specific industry.

6.3 · The pattern to look for

THE ONE QUESTION A PRODUCT BUILDER CAN ASK

"Where is a highly-paid professional doing work that is STILL primarily reading, synthesizing, and filling out forms?"

FIGURE 07 · THE PATTERN

Before and after: inserting the agent into the workflow

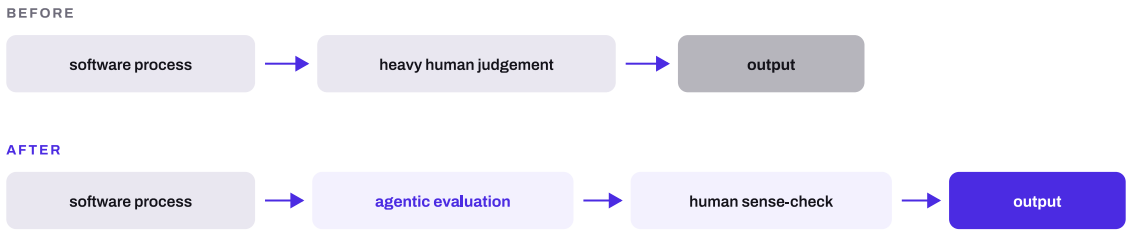

The human moves from the centre of the workflow to a sense-check at the edge. Every instance of that pattern is a candidate for an AI agent.

Every instance of that pattern is a candidate for an AI agent. Every industry where that work happens at scale, construction, legal, healthcare, insurance, finance, is a market where the transition from human-executed cognitive work to agent-executed workflows is inevitable. The question is only who gets there first and with whose data. For a builder with domain depth, e.g. real estate or construction data intelligence, the opportunity is to take what has been built (a passive intelligence platform that showed patterns) and evolve it into an agentic system that acts on those patterns.

PART VII · INDICATORS

What to Watch, Leading Indicators

Fifteen signals in three buckets. Watch the last five as a sign of pace and transformation investment speed.

GO

Indicators the opportunity is real

01 Enterprise AI spending growing 3.2× year-over-year (2024 → 2025).

02 AI deals converting to production at twice the rate of traditional SaaS.

03 50 AI products generating over $100M ARR as of early 2026.

04 Fastest-growing AI startups reaching $5M ARR 1.5× faster than SaaS peers did in 2018.

05 AI-enabled startups commanding 22–86% valuation premiums over non-AI peers.

PAUSE

Indicators that caution is warranted

01 The $500B Sequoia "revenue hole", infrastructure spending exceeds demonstrated revenue generation.

02 SaaS incumbents losing $285B in market cap in a single day as per-seat pricing logic collapses.

03 DeepSeek's cost-disruption suggesting model intelligence will commoditize faster than expected.

04 ROI scrutiny intensifying, enterprises demanding proof of business outcomes, not just demos.

05 Concentration risk: the top 4–5 companies absorbing a third of all global AI venture capital.

WATCH

Key inflection points to monitor

01 Outcome-based pricing adoption rate, speed signals market maturity.

02 Enterprise AI agent production deployments, pilots are everywhere; production is the real signal.

03 Regulatory developments, especially EU and U.S. liability frameworks for agentic AI.

04 Model cost trajectory, if DeepSeek-style compression continues, the model layer becomes fully commoditized.

05 Labor market data, early signs of AI-driven white-collar displacement are the most politically significant inflection.

FIGURE 7.1

The 15-signal quarterly review. ProductMind partners run this with portfolio teams at the start of each quarter.

CONCLUSION

The Strategic Stance

The AI era is real, the revenue is real, and the opportunity is genuine. But the lesson of every major technology transition, from the PC to the internet to mobile, is that the lion's share of the value does not go to those who build the rails. It goes to those who build the trains.

The rails in this era are the foundation models and the GPU infrastructure. They are impressive, important, and largely controlled by well-capitalized incumbents. The trains are the vertical applications, the products built for specific industries, using specific data, solving specific problems for specific buyers whose workflows will need to be upgraded to the current technological level in the coming decades.

For a product thinker with deep domain knowledge in an industry like real estate and construction, where the data model is complex, the workflows are paper-heavy, and the buyers are notoriously skeptical of generic technology, the timing is unusually good. The AI capability is now accessible to drive deeper than the SaaS era could. The domain knowledge required to use it well is rare. The window to establish a position is open, but will not remain open indefinitely.

> THE BUILDER PROFILE
>
> The builders who will matter in this era are not the ones who understand AI the best. They are the ones who understand their domain the best and can translate that into systems that actually do the work. That is a business and PM skill, not just an engineer skill.

THE NEXT STEP

Get specific. Define your beachhead. Understand your data advantage.

The goal of this brief is to give you the conceptual grounding to navigate this landscape confidently, to know which conversations to join, which claims to scrutinize, which opportunities to pursue, and which distractions to ignore. Build something that ships, and learn as quickly as you can along the way.

KEY SOURCES & FURTHER READING

Grounded in primary sources.

Primary references behind the data, analysis, and quoted material in this brief. We have cited the freshest accessible source for each figure; where multiple sources align, we cite the most rigorous methodology.

Menlo Ventures · 2025 State of Generative AI in the Enterprise — Data · Survey

Sequoia Capital · David Cahn, AI Revenue Gap Analysis (2024, 2025) — Analysis · VC

Bain & Company · Will Agentic AI Disrupt SaaS?, Technology Report 2025 — Framework

IDC · Is SaaS Dead? Rethinking the Future of Software in the Age of AI (Dec 2025) — Forecast

Deloitte · SaaS Meets AI Agents, Predictions 2026 — Forecast

Carta · State of SaaS Fundraising Q3 2025 — Data · Funding

Eqvista · AI vs SaaS Valuation Multiples (2026) — Data · Multiples

Gartner · AI Agents Are Disrupting SaaS Pricing (2025) — Forecast

Marc Andreessen · a16z Show AMA, January 7, 2026 — Primary · Audio

ABC News · John Chambers on AI vs Internet Boom (2025) — Primary · Press

Fortune · Eye on AI Newsletter, 2025 Year in Review — Editorial

The American Prospect · "Bubble Trouble", April 2025 — Editorial