



Moneta H1 Market Update

Multiple Compression Is Cyclical, Not Structural

\$285B in software market value erased in 48 hours — part of ~\$2T wiped from the software sector since its October peak — on a fear that was already wrong, and Q1 earnings proved it.

H1 2026

Contents.

1

The SaaSpocalypse

What happened, and why the market got it wrong

2

The Core Thesis

Multiple compression is cyclical, not structural

3

The Mechanism & The Risks

Why AI expands demand — and four risks, steelmanned

4

Capital in Motion

Liquidity is moving, and what it means for LPs

The SaaS apocalypse.

What happened in February 2026 — and why the market got it wrong.

The catalyst.

1

THE FEAR

Anthropic's Claude Cowork launched — plugins automating legal, finance, and marketing: the exact work SaaS vendors charge for per seat.

1 agent ≈ 10 seats

2

THE FICTION

Weeks later, Citrini Research's fictional 2028 macro memo — an AI-driven economic spiral — went viral. Its own preface called it “a scenario, not a prediction.” The market reacted anyway.

9,271 likes · 1,882 reposts

3

THE REPRICING

If one agent does the work of ten licensed users, why pay for ten? In 48 hours, the largest AI-triggered repricing in software history.

\$285B erased in 48 hours

When a fictional memo from 2028 starts moving real money in 2026, sentiment has decoupled from fundamentals.

The selloff in numbers.

\$285B

Software market value erased in 48 hours

22.7x

S&P software forward P/E — first time below the S&P 500, vs. 34x long-run avg

~75%

Of software stocks screening technically oversold at the trough

-55%

HubSpot H1 performance (Jan 1–Jun 30); Monday.com -51%, Adobe -41%, Salesforce -41%

The bellwether declines were indiscriminate.

From the year's start to their lows, SaaS leaders sold off regardless of whether AI was disrupting them or powering them. Goldman's US software basket fell 6% in a single day — its worst since the April 2025 tariff selloff. The iShares Expanded Tech-Software ETF fell more than a third from its September 2025 peak, hitting a 52-week low on April 10.

The market called it the SaaSocalypse; JP Morgan called it “overblown” and “broken logic.” Private markets flinched too — software buyout value fell to \$50B (Jan–May) from \$88B a year earlier.

How it ended.

+33%

Snowflake revenue growth to \$1.39B; AI accounts 9,100 → 13,600 in a quarter

+757%

Dell AI server revenue to \$16.1B, on a \$51.3B committed backlog

+21%

IGV software ETF, off its April 10 low; best month since Oct 2001; -14% YTD

+130%

ServiceNow Now Assist \$1M+ ACV customers, YoY in Q1

Q1 earnings ended the SaaSpocalypse: AI is accelerating enterprise software demand, not displacing it. And private capital followed: H1 2026 set a record for US venture funding.

Snowflake's and Dell's Q1 prints, on consecutive evenings, said the same thing: AI is accelerating enterprise-software and edge-computing demand, not displacing it. The strongest names recovered on the beat and held it — Snowflake ended the half +16% YTD, near its prior peak; Dell hit a fresh high on AI-server demand.

But the recovery split the market — and that's the point. One cohort finished above its pre-selloff highs: at the June 30 close, DigitalOcean was +226% YTD, Datadog +91% (30% above its November 2025 peak), and JFrog +46%. Others bounced on their prints, then re-rated right back down — Salesforce and HubSpot popped 9-10% on clean beats but ended H1 down 41% and 55%. The market re-rated software's risk, not its viability, and is already sorting winners from also-rans. The next bifurcation: which names go lean and AI-native.

The private market agrees — but narrowly.

~90%

of US VC deal value went to AI
in Q1 2026 (\$237B of \$267B)

+56%

AI valuation premium at Series C
vs non-AI (Series B: +55%)

\$19.6M

median Series A round in Q1
2026 — up from \$7.5M in 2020

3.1x


early-stage valuation step-up vs
prior round — the largest since
2015

The capital is real — but the heat is narrow.

AI took ~90% of US VC deal value in Q1 2026, and 73% of it went to just five companies — OpenAI, Anthropic, xAI, Waymo, and Databricks. Beyond the frontier, capital is concentrating in defense tech (\$49.9B across 966 deals in 2025) and cybersecurity — and within AI, applied/vertical software now leads deal activity as the horizontal platforms mature.

The public panic was misplaced — but private capital is narrow: a handful of names, everyone else waiting longer. When capital concentrates this hard, selection is the whole game. That is where our thesis begins.

One market, three signals.



Source: PitchBook-NVCA Venture Monitor (Q1–Q2 2026); US VC Valuations & Returns (Q1 2026); LP Co-Investments in US VC (Q2 2026). US, as of mid-2026.

The Core Thesis.

Multiple compression is cyclical, not structural.

The Perez framework.

Carlota Perez — the economic historian whose *Technological Revolutions and Financial Capital* (2002) mapped how every technological revolution unfolds, from railways to the internet — documented the arc each one follows. The AI repricing fits it precisely.

Installation period

- Capital floods in; valuations detach from fundamentals.
- Speculative excess builds, then a crash follows.
- This is where the SaaSocalypse sits — the installation period of the AI age ending.

The crash

- Not an anomaly — a predictable phase of the cycle.
- The dot-com bust was the internet's crash — the NASDAQ fell 78%.
- A reset of price, not a verdict on the technology.

Deployment period

- Production capital takes over; the tech embeds into the real economy.
- Its deployment era built Google, Amazon, Netflix, and PayPal.
- What comes next for AI. The deployment period is the thesis.

The arc every technological revolution follows



The bust built giants — but not for everyone.

~80%

collapse in US VC investment,
2000 to 2002 (~\$105B to
~\$22B)

\$7.1B

Q4 2001 VC deployed — a third
of the prior year's \$20.9B

2003

fundraising nadir — recovery
then took years

5+

giants launched in the window:
Tesla, SpaceX, LinkedIn, Palantir
founded 2002–03; Salesforce
IPO'd 2004

Capital dried up ~80% — yet the deployment era built a generation of giants. The catch: 2001–03 funds that deployed into the crash were impaired; the winners invested at the nadir and after.

And this is not 2001: the dot-com bust erased ~\$5T over 30 months and took 15 years to reclaim; the SaaSocalypse erased \$285B in software over 48 hours, and recovery is still in progress. The pattern rhymes — the magnitude doesn't.

So there is no cheap-entry trough to wait for, and deployment itself is still a multi-year build. The edge isn't timing the dip — it's selection.

The bifurcation is the signal.

AI-native software now commands a 2–3x revenue-multiple premium over legacy SaaS — roughly 8–15x forward revenue versus 4–6x. The distinction is implementation philosophy — not sector.

AI as an operating system

- Embedded across back-office, product development, and customer workflows.
- Non-technical employees in the product loop; founders with opinions on internal tooling.
- Compounding earnings power, not just multiple expansion.
- Ramp is the archetype — and the playbook is team-dependent, not end-market dependent.

AI as a feature

- Bolted onto the surface and called a roadmap.
- AI lives in the investor update, not in how the team operates.
- Still being priced by the market for displacement.
- Diligence can surface which side a company is on — the pattern is legible early.

Forward revenue multiple — EV / NTM revenue

AI-native software

8–15x

Legacy SaaS

4–6x

Source: FE International, “AI M&A Trends 2026” (Apr 2026) — AI-native ~8–15x vs traditional SaaS ~4–6x forward revenue. Corroborated by Clouded Judgement (Jamin Ball, Jan 2026): public software 12.4x (high-growth) vs 3.1x (low-growth).

Ramp: the proof of concept.

\$7.65B

Ramp valuation, April 2024

\$44B

Ramp valuation, June 2026

60%+

Of Ramp's merged PRs written by its internal agent, Inspect

~6x

Org size vs. 2 years ago — growing as fast as it did then. Meanwhile headcount only grew ~3x

The clearest proof of concept.

Ramp's internal AI operating model runs on three components: Glass (the OS, where one person's breakthrough becomes everyone's baseline), Dojo (a reusable-skills marketplace that extends AI leverage to non-engineers), and Inspect (the coding agent). By May 2026, the typical customer was saving 50% more money and 32% more time YoY.

That is not multiple expansion. It is earnings-power expansion driven by a fundamentally different operating model — and the playbook is team-dependent, not end-market dependent.

Podium: an incumbent reaccelerates.

+300%

AI-agent ARR growth, year on year

\$100M

Podium AI-agent ARR
(Dealroom, Apr 2026)

+30%

customer revenue lift (with
+45% conversion)

10,000+

SMBs live on Podium's 24/7 AI
Employee

Legacy SaaS growth had cooled — then it bet the company on agents.

Podium — an SMB reviews-and-messaging platform whose revenue growth had cooled to ~22% by 2023 — went all-in on its 24/7 “AI Employee.” The agent business crossed \$100M ARR at +300% YoY, with 10,000+ SMBs live and customers reporting ~30% more revenue and 45% higher conversion. Its Jerry 2.0 agent now runs vertical playbooks for automotive, HVAC, and medspa.

Not a new company — the same company, rebuilt around AI. The reacceleration came from the operating model, not the end market.

Fin (fka Intercom): the same move.

+350%

Fin AI-agent ARR growth, year on year

\$3.6B

Salesforce's acquisition of Fin, June 2026

67%

of conversations resolved autonomously (Dec 2025)

\$0.99

per resolution — outcome-based pricing, not seats

A mature support business, rebuilt around an AI agent.

Intercom was a mature, slower-growing customer-service platform — \$382M ARR at year-end 2025 (Sacta); EBITDA not disclosed. It rebuilt around its Fin AI agent, resolving ~67% of conversations autonomously across 40M+ tickets, priced at \$0.99 per resolution rather than per seat. Fin crossed \$100M+ ARR at +350% YoY; the company renamed itself Fin in May 2026, and Salesforce agreed to acquire it for ~\$3.6B in June 2026.

Ramp built its AI-native operating model while still young and growing fast; Podium and Fin prove the same playbook works even after growth has cooled — which is most of a portfolio. That is Enterprise++ applied to incumbents.

One playbook, four moves.

Podium and Fin prove reacceleration isn't only for the AI-native — both were mature and slowing before the pivot. The moves that got them there are concrete, and they apply across the portfolio.

Rebuild the product

- Rebuild around agents — don't bolt AI on. Podium went all-in on its AI Employee (~22% growth → \$100M AI ARR, +300% YoY); Fin rebuilt around its AI agent.
- Run the operating model internally — the Ramp playbook: AI as infrastructure, non-technical teams in the loop, output scaling faster than headcount.
- Neither company waited to be AI-native from birth. The rebuild came first; the reacceleration followed.

Change the model

- Get ahead of pricing: seats → usage → outcomes. Fin priced autonomy directly at \$0.99/resolution — build usage telemetry now, before a buyer forces it at renewal.
- Own what the models can't replicate — proprietary data, workflow, compliance, and distribution.
- Diligence can surface which side of this a company is on — the pattern is legible early, and it is our job to push portfolio companies toward it.

The takeaway: This is not a one-off Ramp story. It is a repeatable playbook, and driving it across every portfolio company — not just the AI-native ones — is Moneta's work in H2 2026.

Our portfolio is already making the turn.

flo

FLO — FKA FLO RECRUIT

Went AI-first company-wide in March 2026 — every engineer on Claude Code, one shared set of practices. Shipping Flo AI: a natural-language layer over their reporting product, built on a semantic layer so answers stay consistent, not just plausible.

40% engineering productivity gain
since going all-in on AI

H Humanly

HUMANLY

Started as SMS/chat screening and interview note-taking, from day one — the data pipeline that let its AI agent launch ahead of peers. In founder Prem Kumar's words: "gasoline, then the car was invented."

Day-one data pipeline → AI avatar
interviews

sikka.ai

SIKKA SOFTWARE

Eighteen years of proprietary, claims-grade data across dentistry, veterinary, optometry, chiropractic, and physician practices — the substrate now powering AI that runs the practice. The moat was built long before the AI.

Eighteen years of data → AI operating
layer

Ramp, Podium, and Fin are the market's proof. These are ours — pre-ChatGPT companies on AI-native trajectories, running the same playbook.

Enterprise++: the thesis, evolved.

Software++ was our founding frame: enterprise software as the compounding core of venture returns. Enterprise++ is its evolution for the deployment period — and it keeps moving. It is the frame we bring to every founder conversation today, not a fixed one.

What Enterprise++ means

- Software alone no longer implies AI — the premium accrues to enterprises that run on it end-to-end.
- The “++” is the operating model: proprietary data compounding into product, non-technical teams in the product loop.
- Moats live in compliance, governance, and workflow ownership — layers the foundation models cannot replicate.
- Beyond software: physical AI, in software-enabled hardware and hardware itself. Proof point — Argon Mechatronics (robotics).

How we put it to work

- Selection: the diligence bar is AI as operating infrastructure, not a feature set — and which side a company is on is legible early.
- Portfolio: institutionalizing the Ramp playbook — Glass, Dojo, and Inspect as organizational principles, regardless of end market.
- Pitches: we tell founders directly — we are Enterprise++, and we help portfolio companies become it.

The one-line version: Software++ said enterprise software would compound. Enterprise++ says the compounding now comes from how the enterprise itself runs. Same conviction, sharper filter — in our pitches, our diligence, and our portfolio work.

The Mechanism.

Why AI expands software demand rather than destroys it.

Mechanism 1: Jevons Paradox.

Efficiency creates new demand — it doesn't just reduce old demand. In 1865, William Stanley Jevons observed that more efficient steam engines burned less coal per unit — yet total coal consumption surged for decades, as cheaper power unlocked industrial uses that were previously uneconomical.

The insight is superlinear, not linear. Cheaper inputs don't just move you along the demand curve — they shift the curve outward by making entirely new categories of consumption viable. The demand-creation effect overwhelms the conservation effect.

The pattern repeats across every general-purpose technology. The spreadsheet didn't eliminate accountants — it created modern finance. The internet didn't hollow out commerce — it created e-commerce, digital advertising, and cloud. BLS projects 15% software-developer growth through 2034, five times the all-occupation rate, naming AI as a demand driver.

Applied to AI: Workflows that once required a dedicated analyst or a six-figure contract are now within reach of a mid-market company with a \$50K budget. Gartner projects \$6.31T in 2026 IT spending, up 13.5% (revised April 2026). That's market expansion, not compression — the binding constraint becomes human oversight capacity.

Mechanism 2: Tasks vs. Jobs.

The fear was a jobs argument: AI replaces workers, workers stop buying seats, SaaS contracts. The labor economics say otherwise.

Tasks, not jobs

- MIT labor economist David Autor's task-based model: automation displaces task bundles, not whole jobs — e.g., losing 3 of 10 tasks to AI, not the whole job.
- Freed capacity is reallocated to higher-value work; new job categories form around the technology.
- The honest caveat: where AI is a near-complete substitute for a role, that demand can stay depressed — and losses are sharp and local while gains unfold slowly.

What the surveys show

- Vanguard (Dec 2025): the ~100 most AI-exposed occupations are outperforming on job growth and real wages.
- EY-Parthenon (Jan 2026): 69% of CEOs expect AI to maintain or increase headcount.
- KPMG (Oct 2025): 92% of CEOs plan to increase headcount even while funding AI.
- CEPR (Feb 2026): 12,000 European firms adopting AI saw productivity rise 4% without cutting headcount.

The aggregate direction of the labor market — and the enterprise software market that serves it — is expansion, not contraction.

Mechanism 3: Vibe-coding won't eat enterprise SaaS.

The claim: AI generates whole apps from prose, so enterprises build their own tools and stop paying SaaS vendors. The threat is real — but bounded.

Why enterprise SaaS holds

- Compliance, security, governance, audit trails, integrations, and trust take years to build and embed.
- Years of proprietary customer data is the most durable moat — and AI increases its value.
- Atlassian (down 65% in the selloff) reported Q3 revenue +32% YoY, beating by ~\$220M and raising guidance.
- Vibe-coding doesn't kill SaaS — it kills weak SaaS.

Where the threat is real

- Retool: 35% of enterprise teams have already replaced a SaaS product with a custom build; 78% expect to build more.
- Displacement hits the commodity end: generic workflow tools, simple dashboards, lightweight reporting.
- The security reality bites: Veracode found serious flaws in 45% of AI coding tests; Tenzai found 69 vulnerabilities across 15 AI-built apps.
- What survives: deep data moats, compliance, and embedded workflows.

Again, the bifurcation: vibe-coding accelerates the separation between defensible and indefensible software — faster than the market expected.

“AI isn't killing SaaS. It is exposing which products were never truly essential in the first place.” — TechFinitive, April 2026

Real Risks, Steelmanned.

Four arguments a thoughtful skeptic would actually make.

Risk 1: Pricing model disruption is real.

The per-seat model is being replaced in real time — and the transition is messy. In a 2025 survey of 240 software and AI companies, seat-based pricing as the primary model fell from 21% to 15% in twelve months, while hybrid usage-and-subscription rose from 27% to 41%. IDC predicts pure seat-based pricing will be obsolete by 2028.

The near-term risk is at renewal. Buyers can treat seat counts as a cost-reduction lever just as vendors still depend on seat expansion for net revenue retention. BCG found 40% of IT buyers name seat reduction as their primary way to cut software spend — softer expansion before visible churn.

Outcome-based pricing is not a clean escape hatch. 47% of buyers struggle to define measurable outcomes; 25% cite difficulty aligning on value attribution. As vendors charge for work completed rather than access granted, investors must ask: scalable software economics, or customized delivery in disguise?

Why the thesis holds: Hybrid pricing is the likely transition state. Companies that meter usage, define defensible value metrics, and show improving unit economics retain software-like valuations. ICONIQ AI builders expect ~52% gross margins in 2026 — below mature SaaS, but the direction matters more than the snapshot.

21% → 15%

Seat-based as primary pricing model, in 12 months

27% → 41%

Hybrid usage + subscription pricing, same period

40%

of IT buyers name seat reduction as their #1 cost lever

by 2028

pure seat-based pricing obsolete (IDC forecast)

Risk 2: Foundation models moving downstream.

Anthropic, OpenAI, and Google are not neutral infrastructure. They are moving up the stack into application-layer territory.

The platform risk is real

- Since Claude Code (Feb 2025), Anthropic has shipped a dozen-plus enterprise connectors (Gmail, Drive, DocuSign, LegalZoom, FactSet, and more).
- Software buyout deal value fell to \$50B in Jan–May 2026, from \$88B a year earlier.
- Anthropic’s \$1.5B enterprise venture (Blackstone, Goldman, Sequoia); OpenAI raising \$4B+ for “The Deployment Company.”
- Model providers are moving toward the implementation work once owned by IT services firms.

Why it’s not existential

- Incumbents own the data, the compliance and governance, and the hybrid human-plus-agent workflows.
- They have to prove it before OpenAI or Anthropic shows it can do the job as well.
- The risk compresses the addressable market and raises the defensibility bar — it doesn’t kill B2B SaaS wholesale.
- Most at risk: thin workflow layers on commodity models with no data moat.

\$88B → \$50B

Software buyout deal value, Jan–May (YoY) — investors pricing the AI uncertainty

12+

enterprise connectors Anthropic has shipped since Feb 2025

Risk 3: Deployment takes longer than expected.

The Perez framework is not a guarantee of timing. Deployment periods are real — but slow, and not automatic.


The Productivity J-Curve

- New general-purpose tech first depresses measured productivity as firms fund complementary change.
- Electrification only paid off once factories were redesigned around it — over decades.
- Brynjolfsson et al.: U.S. TFP was 15.9% higher than official measures once intangibles were counted.
- 88% of firms use AI somewhere; only ~1% consider themselves mature in deployment.

The venture timing problem

- Fund lives are ~10 years; value creation that lags by years compresses IRR.
- The thesis can be directionally right but temporally miscalibrated.
- Mitigation: favor companies where AI is already embedded in product, operating model, and workflow.
- Implementation as a present source of retention and margin — not a future roadmap item.

The Productivity J-Curve — returns lag adoption, then compound



Risk 4: Concentration & the long-tail problem.

If capital and demand concentrate early around a small group of AI-native winners, the long tail of competent SaaS gets re-rated down.

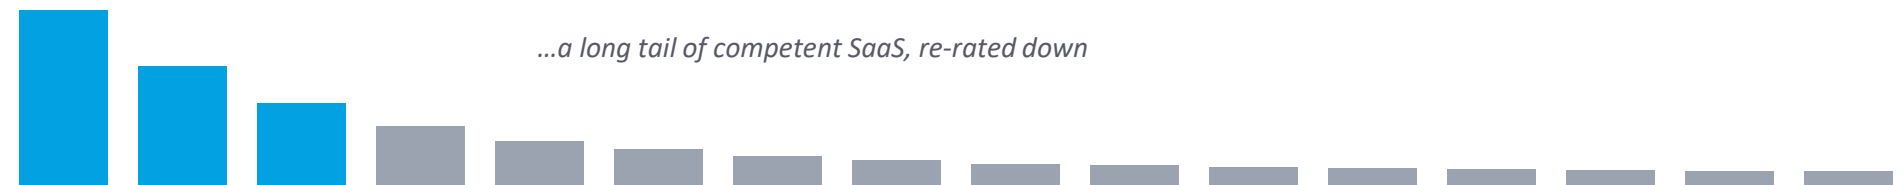
The concentration dynamic

- “Good SaaS business” and “venture-scale winner” become materially different categories.
- Frontier AI economics favor scale: high fixed costs, scale economies, downstream extension.
- One AI-native compounder can outweigh a portfolio of capable point solutions.
- Vertical software’s workflow ownership and data are necessary but no longer sufficient.

How early-stage is positioned

- Adjustment costs are hardest for older firms; younger, growth-oriented firms adapt better.
- Startups carry less legacy workflow and organizational inertia.
- But the bar for selection rises: AI must become operating infrastructure, not a feature set.
- Moneta’s Enterprise++ thesis (our evolution of Software++) is positioned for the bifurcation — not insulated from it.

A few AI-native winners



The concentration curve — a few AI-native winners take the capital; the long tail gets re-rated down.

Capital is in motion.

\$1T

Anthropic IPO filing (June 1),
targeting October

\$75B

SpaceX IPO raise (June 11) —
the largest IPO in history

\$60B

SpaceX's all-stock acquisition of
Cursor, four days later

\$66B

Cerebras first-day valuation
(mid-May), priced at \$185

Four headline events in the first two weeks of June.

Anthropic and OpenAI both filed to IPO; SpaceX went public in the largest offering ever, then immediately spent newly minted stock acquiring Cursor; Cerebras had already opened the window. This is the public-market half of the cyclical-not-structural thesis confirming itself in real time.

SpaceX redeploying freshly minted liquidity into an acquisition rather than banking it is the clearest data point that capital does not sit still once it is freed.

What this means for LPs.

Liquidity is moving again. Distributions from public holdings, IPO proceeds, and secondary exits are creating redeployment capital at a scale the market hasn't seen in years. SpaceX immediately recycling IPO stock into the Cursor acquisition is the clearest sign that capital does not sit still once freed.

Early-stage venture is positioned to capture a disproportionate share. Entry valuations at seed and Series A are negotiated deal-by-deal and not subject to the narrative momentum that inflated public comparables in 2024–2025. Where multiples have reset and outcomes will be dispersed, the early-stage entry point offers a margin of safety that mid-to-late-stage exposure does not.

The Perez framework becomes a fund-timing thesis. The deployment period is when AI stops being a growth story and becomes infrastructure assumed by every enterprise buyer. The companies that define it are being built now — and in many cases are already in our portfolio.

This vintage will be defined by the funds that tell the teams building Glass from those merely cleaning up their investor updates — and by helping our founders be the former. That is where our energy is focused in H2 2026.

What this means for founders

The repricing told us exactly what the market pays a premium for. Wherever you sit, the moves are concrete — and H2 2026 is the window to make them.

If you sell software today

- Move AI from the roadmap into the operating system — the 2-3x multiple gap separates companies that run on AI from companies that mention it.
- Get ahead of pricing: seats → usage → outcomes. Build usage telemetry now, before a buyer forces it at renewal.
- Take the vibe-coding test honestly: if your product could be rebuilt in a quarter, the moat is the data, workflow, and compliance around it — invest there.

If you are AI-native

- The 2-3x premium and 8-15x multiple is earned quarterly — it reprices fast if the operating model does not show up in the numbers.
- Run the Ramp playbook internally: one person's breakthrough becomes everyone's tool — 6x the valuation on 3x the headcount.
- Foundation models are moving downstream. Own what they cannot replicate: your data, your workflow, your distribution.

H2 2026: What this means for Moneta

With Fund III substantially deployed, the work of H2 2026 is not primarily about what we invest in next. It is about what we build inside the companies we already own. In a deployment-period market, portfolio support is the product.

Defense: protecting the portfolio

- Navigate the pricing transition — build the telemetry and value metrics to defend NRR as buyers treat seats as a cost lever.
- Reinforce the data moat — turn years of proprietary customer data into product advantages that can't be replicated.
- Honest bifurcation assessment — have the uncomfortable defensible-vs-indefensible conversation now, not at the next raise.

Offense: driving AI adoption

- Drive real internal adoption – AI changing how the team works day to day, not just what the product does
- Get non-technical people into the product loop – adoption compounds when the whole team builds with it, not just engineering
- Make internal adoption a standing agenda item – use the four questions below as a shared diagnostic for where each company stands

The four questions: (1) Where is AI already changing how the team works, not just what the product does? (2) Who on the non-technical team is in the product loop? (3) Is the proprietary data advantage compounding into product? (4) Can we defend NRR at the next renewal with real telemetry and value metrics?



Thank you. ■

www.moneta.vc