

— STRATEGIC INSIGHTS

A Framework for B2B Software Valuations in the AI Era

*What still drives value, what AI changes, and how
to decide what to do about it*

CONTENTS

Table of Contents

—	Summary	3
01	The Software Market Reset: A Postmortem	5
02	Who This Report Is For: The Four Company Archetypes	6
03	The AI Defensibility Overlay Analysis	7
04	Core Valuation Drivers: What AI Changes (and Doesn't)	8
05	Three New Diligence Issues Specifically for AI	14
06	The Coming Pricing-Model Transition	16
07	The Buyer Side: Who Pays for What	17
08	Transaction Timing: Invest, Hold, or Transact	19
09	AI Defensibility Case Study: Grow Faster vs. Grow Stronger	20
10	The 85 Advisors Forward View	22
11	What We're Watching, and What Could Change	23
—	Appendix: AI Defensibility Diagnostic	24
—	Definitions	25
—	Sources	26

SUMMARY

Executive Brief

Two shifts in valuation, five findings, and three actions for your pre-process checklist

Two things have happened to software valuations since 2021, and they are routinely confused with one another. The first is a market-wide repricing, driven mostly by interest rate increases and by a reversion to the mean (a fall back toward historical averages) from an unsustainable peak. That affected the entire asset class. The second is a shift in how buyers judge the durability of SaaS revenue in the face of AI, and that judgment is made company by company.

Five Findings

AI defensibility now sets the confidence interval around the standard software valuation metrics. The valuation models used by acquirers are still underpinned by company growth, retention, and scale. AI layers a question on top of the numbers: do a company's customer dependence, switching costs, data position, and unit economics make its current performance more durable, or less? The metrics say what the business earns today, while defensibility focuses on how much of that performance a buyer believes will still be there years from now. Buyers are pricing and structuring deals on that belief.

ARR growth and NRR are the two dominant valuation levers. Published models bear this out. For example, SaaS Capital's 2026 software valuation model prices a company off just three inputs: the public market SaaS revenue multiple, growth in ARR (annual recurring revenue), and NRR (net revenue retention, or what last year's customers spend this year, expansion net of churn). Within that model, growth carries roughly twice the weight of retention. In our experience, buyers weigh NRR more heavily than the arithmetic suggests, and pay a premium for it. High NRR is the clearest proof of the durability they are trying to price. It shows expansion within an existing customer base, and switching costs high enough to hold it there. The higher and more durable the NRR, the further the buyer's thesis moves from "good business" toward "self-funding growth machine."

Gross-margin benchmarks are becoming business-model-specific, not lower across the board. A sub-70% gross margin can survive M&A due diligence, but only with strong proof behind it. Buyers want to see three things: AI features that are central to the product's differentiation, per-unit costs falling with scale, and documented offsetting cost savings. Absent those proofs, a lower gross margin is simply a discount with a narrative attached.

Different buyers look for and pay for different things. A strategic acquirer, a PE platform, a PE add-on, and a growth-equity or AI-native buyer (each defined in Section 07) may value the same company on different drivers. The same company, with the same metrics, is worth more to some of these buyers than to others, and the spread between your best and second-best offer may come down to which of them are competing for you.

M&A timing is company-specific, not a market-timing question. Many sellers ask which calendar window will be most receptive to their M&A process. However, the question that matters more is whether your evidence of durability in the face of AI threats is strengthening faster than your competitive risk. Section 08 provides observable triggers for making that judgment.

SUMMARY

Three Actions for Your Pre-Process Checklist

ACTION 01**Run the AI defensibility diagnostic in the Appendix**

Identify how your business scores on AI defensibility dimensions and decide whether potential weaknesses are structural or addressable before initiating an M&A process.

ACTION 02**Start building the evidence file to support the KPIs buyers will ask for**

NRR recalculated from customer-level data, AI unit economics with a documented cost trajectory, a data-rights file, and a considered answer on post-close operational AI initiatives that could drive efficiencies.

ACTION 03**Map your realistic buyer universe before settling on process timing**

The invest, hold, or transact decision in Section 08 also turns on the profiles of realistic buyers, as well as what that class of buyer is looking for and typically pays for.

SECTION 01

The Software Market Reset: A Postmortem

Four causes of the multi-year valuation compression, and why AI takes more blame than warranted

It's tempting to read the multi-year decline in software multiples as a verdict on AI. It isn't, mostly. We see four distinct causes: three are market-wide and outside any company's control. The fourth, AI defensibility, is something a company can actively manage, and it's a significant focus of the rest of this report.

Rates did the heavy lifting first. Software is a long-duration asset. Most of its value sits in cash flows years away, so when interest rates rose off their floor, the present value of those distant cash flows compressed harder than almost any other equity category. The relationship is well documented: when the five-year Treasury yield rises, public software multiples fall. The first, largest leg of the valuation compression was a rate story.

The software asset class was overvalued to begin with. The lofty valuation multiples of 2020–2021 were an artifact of near-zero interest rates, pricing in a durability of growth that few businesses could ever deliver. A reversion was coming with or without AI. PitchBook analysis attributes part of the market's slow adjustment through 2024 and 2025 to anchoring: investors still measuring valuations against the extraordinary conditions of the pandemic peak. Some of what gets called a collapse is simply air leaving a number that was never going to hold regardless.

AI became the narrative that completed the repricing. In our view, by the time AI-disruption fear became the dominant story, software multiples were already stretched against history and already carrying an interest rate headwind. Markets often need a coherent and convenient story to finish a repricing they were primed for anyway, and AI supplied it. AI risk is not imaginary; it is precisely what the rest of this report is about. But AI has been doing double duty as both a genuine driver of valuation risk for exposed companies and a convenient justification for an overdue markdown. The most recent leg down in valuations came even as rates eased, and was widely attributed to AI risk. We interpret this as the market settling on AI as the explanation for a discount it was inclined to take anyway.

The broader risk environment repriced. A less predictable macro and geopolitical backdrop appears to have raised the equity risk premium (the extra return investors demand for holding stocks) generally, software included. A company should not attribute to AI a discount that partly reflects the wider world.

The practical takeaway: Most of what happened to software multiples over the last three years was done to the asset class: by rates, mean reversion, and macro factors. This report examines the thing that is in a company's control: its AI defensibility position.

SECTION 02

Who This Report Is For: The Four Company Archetypes

Find your archetype before reading the drivers

85 Advisors typically serves B2B software companies with revenue between \$3M and \$30M. Within that range, the AI defensibility dimensions defined in Section 03 play out differently by company archetype, as do the buyer universes each archetype attracts, which is covered in detail in Section 07. Readers should anchor themselves accordingly.

ARCHETYPE	TYPICAL DEFENSIBILITY PROFILE	AI-ERA PRESSURE POINT
Vertical SaaS with regulatory or compliance depth	High system-of-record status; high non-software complements; senior, high-stakes users	Showing that AI features strengthen the core product's defensibility, not just add a feature on top
Horizontal workflow / productivity tools	Feature-led differentiation; moderate switching costs; broad but shallower usage	Commoditization of "good enough" functionality by AI-native entrants
Data / intelligence businesses	Value concentrated in the data asset itself; thin services layer	Substitution where the underlying data can be sourced or synthesized elsewhere
AI-embedded / usage-priced products	Differentiation via model + workflow; COGS carries inference	Margin structure, revenue-model comparability, and pricing-transition volatility

Archetypes are simplifications, and many companies blend two.

SECTION 03

The AI Defensibility Overlay Analysis

Four defensibility dimensions measured on two axes — AI exposure vs. AI readiness

AI has moved defensibility to the center of due diligence. We use a four-part heuristic to assess it. We borrow three factors from SaaS Capital's AI Assessment Framework, while the fourth is our own addition.

DIMENSION	WHAT IT MEASURES	SOURCE
System of Record (SoR)	Whether the product houses the definitive “ground truth” of a customer's operations, or merely analyzes and processes data that lives elsewhere	SaaS Capital
Non-Software Complements (NSC)	Value beyond pure code: services, compliance, hardware, marketplace, proprietary data access: things a generalist AI cannot simply regenerate	SaaS Capital
User & Usage (U&U)	The seniority of the users and the monetary stakes of the decisions or transactions the product touches. High-contract-value products facilitating big-dollar, executive-level processes are structurally harder to displace, and executives cut tools for the lower ranks before the ones they personally use	SaaS Capital
Workflow Embeddedness	How deeply the product is woven into daily operating rhythm: frequency of use, breadth across teams, and the re-implementation, retraining, and compliance rework a customer would incur to leave	85 Advisors

Beyond Disruptor vs. Disruptee

Most software businesses both use AI and face AI substitution, in different layers of the same product. There are two axes, and the direction that a company is moving on each axis is a more important signal than its current position.

AI displacement exposure asks how much of the product's value sits in layers a general-purpose AI model can increasingly approximate. Low system-of-record status, low non-software complements, low-stakes usage, and shallow embeddedness taken together define the vulnerable company profile.

AI value-capture readiness asks how well positioned the company is to monetize the AI shift: its data position, its architecture, its unit economics, and its ability to sell the automation of work rather than assistance with it. PitchBook's 2026 enterprise-AI research evaluates readiness across eight frameworks, including data-moat depth and agent-operations maturity, and we track similar questions in diligence preparation.

SECTION 04

Core Valuation Drivers: What AI Changes (and Doesn't)

Our revised valuation framework of eighteen drivers: updated to show what AI changes about each — if anything

We built this framework to help business owners and leaders understand the factors buyers use to set valuation, and provide a roadmap for maximizing enterprise value ahead of a potential sale or financing. A revenue multiple driven by growth and retention is a useful anchor precisely because it's simple, but disaggregating those inputs reveals what sits behind the number: the drivers an owner can actually influence. This framework builds upon our first edition of valuation guidance, updating it for today's AI context.

Acquirers do not buy software companies; they buy cash-flow streams: discounted, risk-adjusted, and modeled forward. Each driver we discuss maps to three essential questions:

- How predictable are the revenue and cash flow?
- How defensible are they?
- What will those revenue and cash flow streams look like in the future?

Although this logic still holds, AI has changed the context. For each valuation driver, we describe what AI has changed in either the interpretation of the metric or the metric itself. We have also layered the concept of a company's AI defensibility over the entire analysis. At the same time we simplified the model, collapsing our earlier four-tier stratification into three clearer bands (minimum, strong, and premium) to provide a more useful resource that addresses the questions buyers are asking today.

Each group opens with its reference table, a single set of tier thresholds per metric used consistently throughout the report, followed by driver-by-driver commentary. Our thresholds and weightings are directional judgments rather than a precise valuation or scorecard. We encourage readers to use this guidance to locate where your business sits on each driver and where to concentrate before initiating a process.

SECTION 04

Quantitative Software Valuation Drivers

DRIVER	— TIERS —			WEIGHT	WHAT AI CHANGES
	MINIMUM	STRONG	PREMIUM		
1. Revenue scale (ARR)	\$3M+	\$10M+	\$20M+	Low	Interpretation. Thresholds hold; defensible companies draw interest earlier in the revenue range.
2. Recurring revenue %	70%+	80%+	90%+	Medium	Metric. AI consumption revenue complicates ARR definition and comparability.
3. ARR growth	10%+	20%+	30%+	High	Interpretation. Compare against AI-native rivals. Whether rising or falling matters more than the level.
4. Gross margin (subscription revenue)	65%+	70%+	80%+	Medium	Metric. For AI-embedded products, the benchmark depends on whether compute cost is passed to customers or carried in margin; proofs required.
5. EBITDA margin	0%+	10%+	20%+	Low	No change. Reinvestment-for-defensibility case strengthens. Independent of AI, profitability always matters to some (and to all when capital is expensive).
6. Cust. Concentration (largest vs. top 10)	<20% <60%	<15% <50%	<10% <40%	Medium	No change. Hyperscaler concentration carries both strategic value and dependency risk. Where AI defensibility is unproven, concentration risk weighs heavier.
7. Rule of 40	20+	30+	40+	High	Interpretation. Harder to clear, because AI raises the cost of growth and the cost of the product at once. Companies still above 40 are either pricing AI into the base or offsetting it with internal savings.
8. GRR	85%+	88%+	90%+	High	Interpretation. Scrutinized closely to diagnose structural churn from AI-native substitution.
9. NRR	100%+	105%+	110%+	High	Weighting. Thresholds hold; NRR and growth are the two dominant levers. High NRR implies AI-native rivals cannot displace you: customers expand vs. leave.
10. LTV:CAC	3:1+	4:1+	5:1+	High	Interpretation. Rising acquisition costs and improving retention pull in opposite directions. What matters is your definition and your payback period.

Tiers read left to right as minimum targeted level, strong, and premium. Gross-margin tiers shown are for subscription models; AI-embedded products are judged on per-unit economics and trajectory, not a single range. EBITDA tiers are stage-dependent. LTV:CAC assumes CAC payback under 18 months. NRR below 100% indicates a contracting base but is not, by itself, evidence of AI-driven distress; diagnose the cause. All retention figures TTM, customer-level.

SECTION 04

Quantitative Software Valuation Drivers — Commentary

1. Revenue scale — unchanged. The mechanics of deal sizing hold. Below roughly \$3M of ARR, a company is considered a project to most acquirers; from \$3M to \$10M it is a more viable target; above \$10M it supports premium multiples and genuine competitive tension. AI-defensible companies in the \$5–10M range are drawing buyer attention earlier than they used to, as acquirers try to get ahead of markets they expect to consolidate.

2. Recurring revenue — amplified, but with a new definitional complication. The valuation premium for predictable ARR holds. That said, AI adds a consequential bookkeeping question: features sold on consumption create hybrid revenue that does not fit traditional ARR accounting. Settle how you classify it, and whether your figures line up with the comparable companies (the “comps”) a buyer will benchmark against, before a process begins. See Section 06 for further discussion of the AI pricing-model transition.

3. ARR growth — amplified, and still the largest company-specific lever. The thresholds hold (20% and above is strong, 30% and above is premium; see the tiers above), but the reference point changes. A company growing 15% while AI-native competitors in the same space post 40% is valued differently than it would have been three years ago, not because 15% became bad, but because the environment reset the benchmark. The velocity of travel relative to those entrants now carries as much weight as the level itself.

4. Gross margin — thresholds becoming business-model-specific. The traditional 75–80% gold-standard benchmark was built for businesses with minimal cost of goods. For products carrying material AI inference cost (the computing bill incurred each time the AI runs), some buyers now accept lower margins, but no published dataset establishes a new “acceptable range,” and “AI-intensive” covers business models far too different for a single number to serve them all. In our observation, the companies earning premium multiples despite sub-70% margins arrive with three proofs: their AI features are central to differentiation rather than cosmetic; their per-unit compute cost is declining as the product scales; and their offsetting operational savings are documented, landing in EBITDA typically three to twelve months after the COGS hit. In your own accounts, it is preferable to separate inference cost from implementation cost from human-delivered work, and evaluate contribution margin per unit of usage, pricing coverage, and vendor concentration rather than the consolidated margin alone.

5. EBITDA margin — unchanged. The split between ARR-multiple buyers and EBITDA-multiple buyers holds, and so does what each optimizes for. AI strengthens the case for reinvesting margin into defensibility at the growth stage, provided the spend produces measurable competitive advantage rather than simply consuming margin.

SECTION 04

Quantitative Software Valuation Drivers — Commentary, Continued

6. Customer concentration — unchanged, and best tracked as two metrics. Measure the largest customer and top-ten concentration separately, with ARR as the denominator. Aim for the largest customer to be under 10% of ARR, as anything over 20% results in a material discount (the table above has the full tiers). A concentrated relationship with a hyperscaler cuts both ways. To some buyers it is a strategic asset; others will focus on the platform, renegotiation, and change-of-control risks.

7. Rule of 40 — amplified, and possibly harder to clear. Growth rate plus EBITDA margin above 40 remains the cleanest single measure of efficiency. However, AI is making this threshold harder to reach: it raises the cost of growth through go-to-market intensity and the cost of product through compute in COGS. Companies holding above 40 against those pressures are either pricing an AI-enhanced product into the existing base, or taking enough internal cost out with AI to offset the product-side increase. Buyers will notice this discipline.

8. Gross revenue retention — unchanged threshold, new diagnostic. The 85% minimum holds, with 90% and above as the premium tier. However, a weak number may now signal customers are leaving for AI-native alternatives that deliver most of the functional value at a fraction of the cost. This type of churn requires deeper product or business-model repositioning than an ordinary roadmap contemplates. Ordinary causes of churn remain: service quality, customer health, pricing, implementation, segment mix, and macro conditions. Our AI defensibility framework is a useful tool for distinguishing between types of churn, mapping how displacement churn rises where system-of-record status, complements, stakes, and embeddedness are all low.

9. Net revenue retention — elevated. High NRR demonstrates an incumbent's defensibility relative to new entrants: revenue expansion from existing customers, high switching costs, and compounding relationship value. This meaningfully impacts a buyer's offer: the higher and more durable the number, the further the thesis shifts from "good business" toward "self-funding growth machine," and premium valuations follow. Today, the highest-return NRR lever for companies in our range is likely AI functionality introduced as priced expansion (automation modules, analytics tiers, copilots) built to fit the intended workflow from the start.

10. LTV:CAC — emphasized, with both sides of the ratio moving. Customer acquisition costs are rising where AI-native players are buying share; lifetime value improves where AI-driven retention and expansion are working, so the net effect is company-specific. The 3:1 minimum holds as a floor. Successful sellers define the formula, show CAC payback (the months needed to earn back the cost of winning a customer) alongside the ratio, and can explain the story behind the number, including, if the ratio is very high, why it reflects efficiency rather than underinvestment in growth.

SECTION 04

Qualitative Software Valuation Drivers

DRIVER	RANGE		WEIGHT	WHAT AI CHANGES
	WEAKEST	STRONGEST		
1. Delivery model	License + services	Cloud / modular	High	Interpretation. Gap widens; modern architecture and governed data generally required to ship AI rapidly.
2. Product differentiation	Commoditized	Proprietary logic, compliance, network effects	High	Metric. Products competing only on basic features are worth less; AI makes those features cheap and easy to replicate.
3. Technology / architecture	Legacy	Modern, AI-ready	High	Metric. Is your architecture something a buyer can extend with AI, or something they would need to rebuild?
4. Market growth	<5%	15%+	Medium	Interpretation. New question: is AI expanding the niche or commoditizing it?
5. Competitive dynamics	Unfavorable	Clear niche leadership	High (was Medium)	Weighting. As AI clears out weaker players, being the leader in your niche carries more value than before.
6. Management team	Founder-dependent	Institutional depth	Medium	Metric. AI literacy added to evaluation: can leadership articulate what AI can and cannot do for the roadmap?
7. TAM	<\$200M	\$1B+	Low*	Metric. Products that eliminate work can claim labor budget as TAM. Show the bridge: which roles, what share of loaded cost, and who owns that budget.
8. Market position	Negligible	Recognized leader	Medium (was Low)	Weighting. Leader holds the data, relationships, and embedded usage hardest to replicate.

*TAM warrants materially higher weight only where the product demonstrably eliminates work rather than assisting it.

SECTION 04

Qualitative Software Valuation Drivers — Commentary

1. Delivery model — reinforced. The cloud delivery model premium holds and has widened, because modern, modular architecture with governed, accessible data generally ships AI features at competitive pace, whereas legacy delivery models generally cannot. Buyers focus on architecture and data readiness rather than hosting doctrine, as in some circumstances, private-cloud and on-premise AI deployment can be entirely appropriate, particularly in regulated verticals.

2. Product differentiation — repriced at the bottom. AI has made “good enough” dramatically cheaper to build. While the top tier (proprietary logic, embedded compliance, network effects, high complements) remains relatively unchanged, products competing on features face commoditization pressure from AI-native point solutions. As a result, previously mid-tier products (those that are useful, well executed, but not unique) now score closer to the bottom if AI can approximate their core function. Buyers consider how fast a competitor could replicate your product's features, focusing on architecture review and the data-moat tests described in Section 05.

3. Technology and architecture — emphasized, with readiness added. Buyers still evaluate the architecture that predicts maintenance cost and feature velocity. They now also assess how quickly they could deploy AI on your stack after close, potentially probing agent-operations maturity explicitly. If acquirers feel they are buying a rebuild rather than a platform, valuation is negatively impacted.

4. Management team — unchanged, with one new qualifier. Buyers still primarily test whether the business scales beyond the founder. They now also consider AI literacy and whether the CEO and product leadership can articulate what AI can and cannot do for this specific roadmap.

5. Total Addressable Market (TAM) — the labor-budget reframing, with the math to back it up. The emerging “service-as-software” thesis holds that AI lets software compete for labor budget rather than software budget, selling completed workflows instead of per-seat tools, against the multi-trillion-dollar knowledge-work payroll. Our view is to both appreciate the argument and caution against overclaiming it. Payroll does not convert dollar-for-dollar into software spend: procurement ownership shifts, risk and implementation absorb value, and customers demand a share of the gains. Sellers should present labor budget as a TAM expansion case with explicit assumptions about revenue conversion (including which roles, what share of the loaded cost (salary plus benefits and overhead) is realistically capturable, and who owns that budget) if the product genuinely eliminates work rather than assisting it.

6–8. Market growth, competitive dynamics, and market position. A growing market still earns a premium; the new question is whether AI is expanding your niche or quietly commoditizing it, and whether your roadmap is driving the expansion or merely defending the base. Meanwhile, AI accelerates consolidation in fragmented markets; the threat bearing down on you is bearing down on every other incumbent too, so the acquirer thesis increasingly becomes “buy the leader before AI clears the field.” That raises the weight on clear niche leadership, and on the data, relationships, and embedded usage the leader uniquely holds.

SECTION 05

Three New Diligence Issues Specifically for AI

Three topics that were not on many diligence checklists three years ago. The fourth item is a lens on the buyer; discussed in more detail in Section 07.

ITEM	WHAT IT CAPTURES	WEIGHT
1. Proprietary data / data moat	Rights and governance (can you legally use the data?); dynamic vs. static data (does it refresh?); performance dependency (does the product get better with it?); replication difficulty (could a rival recreate it without your installed base?).	High
2. AI exposure & readiness (two axes)	How much of your business could AI displace, and how ready are you to capture value from it? Both can be inferred from the drivers above; which way you are moving matters more than where you sit today.	High (overlay)
3. Compute cost structure	What it costs you every time the AI runs. Buyers want to see these inference costs as a line in COGS, priced per unit of usage, with evidence the cost is falling as you scale.	Medium (High if AI-embedded)
4. Buyer-side: PE operational diligence	A buyer's own plan for AI-driven improvement post-close. It shapes their due diligence questions, but not necessarily the headline valuation unless realized.	Buyer-specific

1. The Data Moat: How You Know You Have One

AI performance scales with proprietary data, and the depth of a data moat has become a standard diligence question. The distinction that decides whether data is a moat is static versus dynamic. A one-time dataset, with no feedback from usage, is something a funded competitor can eventually assemble. A moat exists when performance improves as customers use the product, compounding with the installed base: more customers yield better models, which create more value, which attract more customers. Once that loop is established, it is very hard to reproduce without the customer base itself. Three tests tell you whether you have one:

Is the data dynamic or static? If the product is equally capable on day one and after three years of customer activity, data is not your defensibility driver.

Does performance measurably degrade without the proprietary data? If a general-purpose model on public data gets close to the same outcome for the tasks customers actually pay for, the moat is thin. Show the delta rather than describing it, and judge “close” task by task, on error cost, human-review burden, and switching economics.

Would replication require your customer base, or just compute time? True data network effects require the installed base. If a rival can match it with one training run, it is not a moat.

Proprietary data is a moat only if it is legally usable, with contractual rights to train on customer data, privacy compliance, customer consent, and the ability to combine data from multiple customers into a single training set rather than keeping each customer's data walled off. A dataset that a company cannot lawfully use, or that customers can withdraw, fails in diligence regardless of its quality. Buyers increasingly read the data-processing agreements before they read the model benchmarks.

SECTION 05

A verified data moat changes which buyers compete and what investment case they are prepared to put money behind. The valuation effect is specific to the situation. We have seen genuine moats materially strengthen outcomes, but we have no transaction sample supporting a numeric range, so we do not publish a multiple premium for one.

2. Reading the Two Axes in Practice

Two questions: what part of your business could AI take? and how well placed are you to profit from AI? You are more exposed if customers could unplug you without much disruption (you are not the system of record, few other tools depend on you, the cost of a mistake is low, and you are not woven into daily work) and more exposed still if the customers you lose are leaving for AI-native products. You are more insulated if you hold usable proprietary data, your architecture can ship AI quickly, your unit economics hold up, and you sell results rather than seats. Most companies sit in the middle of both. For now, buyers care most about which direction you have been moving over the past year, though we expect buyer expectations to continue to rise over time.

3. Compute Cost Structure

For products with material inference cost, compute is now a recurring COGS line with its own unit economics, and buyers treat it that way. The dynamic mirrors the gross-margin discussion: costs hit COGS immediately, while offsetting operational savings reach EBITDA months later, and buyers model both. Expect four questions:

- What is AI COGS as a share of ARR, tracked separately from ordinary infrastructure?
- Is cost per unit of usage declining, at what rate, and is the trajectory documented?
- Is the product priced to cover its compute, or is the company subsidizing the experience at the expense of margin?
- What does the margin profile look like at twice today's ARR under a conservative cost trajectory?

Although buyers do not expect perfect forecasts, they want to see that a company has taken unit economics seriously enough to have a model.

SECTION 06

The Coming Pricing-Model Transition

From seats to outcomes: the repricing of software revenue

Although the AI defensibility narrative dominates the conversation today, the change in how software gets priced may be the bigger story in the long run. PitchBook projects the software industry's transition from seat-based to consumption- and outcome-based pricing running from roughly 2026 through 2031: vendors selling completed workflows rather than per-seat licenses, with enterprise value shifting from the volume of licenses sold to the quality of digital work performed. The same research is candid about what the shift costs incumbents: trading the predictability of per-seat subscriptions for revenue that rises and falls with delivered outcomes has already unsettled investors, and near-term performance can be muted while vendors reprice. The transition is most reflected in three metrics, each covered elsewhere in this framework:

ARR comparability. An "ARR multiple" on a business that is 40% consumption revenue is not directly comparable with one on a pure subscription base. Report contracted ARR (subscription base plus guaranteed minimums) separately from non-guaranteed usage revenue, and track inference cost alongside hosting COGS so you can see gross margin by customer tier as usage fluctuates. Define what is in your ARR, and know how the comps in your process define theirs.

Retention mechanics. Usage-based NRR bounces around more: expansion and contraction both flow through it faster, so a single quarter can mislead. We advise presenting usage trends by customer cohort alongside the headline number, so ordinary fluctuation is not mistaken for churn.

Benchmark validity. The SaaS Capital valuation model was estimated on subscription businesses. The further your revenue model moves from that subscription-based population, the wider the error on any formula-based estimate, and one more reason the formula is a starting point rather than a price. This has implications for transaction timing. Barring other circumstances, a company actively renegotiating its pricing model with customers may capture a cleaner multiple by selling before the transition than in the middle of it, as revenue recognized cleanly under the old structure (pure subscription-based) shows up differently under the new one (blended subscription and usage-based model). If a pivot is on your roadmap, assess its impact on transaction timing explicitly. How buyers value hybrid revenue is evolving, and calibrating to current buyer practice is part of our standard pre-process work.

SECTION 07

The Buyer Side: Who Pays for What

Different buyers pay for different things

The same company could be valued on different drivers, and often at different levels, by each type of acquirer. The invest, hold, or transact decision in Section 08 depends on which classes are realistic for a specific company.

BUYER CLASS	WHAT THEY PAY FOR	AI-ERA SHIFT	IMPLICATION FOR SELLERS
Strategic acquirer	Product, data, roadmap, and customer access that accelerate their own thesis	Paying up for data assets, embedded workflows, and "buy the leader before AI clears the field" consolidation	Strategic value is buyer-specific; competitive tension is what converts it into price
PE platform	A durable earnings base they can grow and add onto	Testing post-close AI-driven margin expansion as a value-creation plan: probed hard, rarely paid for up front	Understand the P&L-transformation thesis (below) and prepare the operational evidence for it
PE add-on	Synergy into an existing platform: customers, geography, product tuck-in	Platforms moving earlier on tuck-ins as AI thins the field of independents	Add-on math tolerates smaller scale but pays for fit, not standalone potential
Growth equity / AI-native	Readiness to capture the shift: data position, architecture, outcome-based motion	New class of interest in distribution and proprietary data for AI products	Readiness evidence (Section 05) is the diligence file this buyer reads first

The map is for deciding which evidence file to build and which process design creates tension.

In practice most processes attract a mix of these classes of buyers. The work, then, is narrowing to the two or three classes a company can realistically attract, building the evidence file each of them reads first, and designing a process that puts at least two of them in competition at the same time.

SECTION 07

The PE Post-Close Playbook and Hypothetical Return Scenario

With the valuation multiple tailwind gone, PE return models lean on earnings growth planned for after close through systematic AI deployment: back office and finance first, then customer operations, then sales and pricing. The intent is well documented. PitchBook's survey of major PE firms' recent earnings calls found AI treated as “both a risk factor and an operational lever,” central to underwriting discussions. Yet post-close improvement programs have a long history of underdelivering. Buyers know this, which is why the operational-AI opportunity is a subject of diligence rather than something paid for explicitly in the headline price. Where it is paid for at all, it is through earn-outs and similar contingent payments that arrive only if the improvements are realized. Two things follow for sellers. First, buyers pay for upside they can verify. What converts their upside into your price, or into better terms, is evidence they can check, a fit specific to their own plans, and a second bidder at the table; describing the opportunity does nothing on its own. Second, the downside is priced right away. Fragmented data, poor operational visibility, and a messy tech stack come straight off the price, because the buyer has to fix them before any improvement program can start.

Here is an illustrative bridge:

ASSUMPTION	ENTRY	EXIT (YEAR 4)
ARR (\$10M growing ~20%/yr)	\$10.0M	\$20.7M
EBITDA margin (ops program: 15% → 30%)	15% (\$1.5M)	30% (\$6.2M)
Multiple paid / achieved	4.0x ARR = \$40M (26.7x EBITDA)	3.5x ARR = \$72.5M (11.7x EBITDA)
Financing	40% debt (\$16M), \$24M equity	~\$10M debt after paydown
Estimated equity outcome	—	\$62.5M → 2.6x the equity invested (MOIC), 27% annualized return (IRR)

Illustrative only. Note what the bridge actually shows: the return requires sustained growth, a fifteen-point margin program (an assumption many investments fail to achieve), leverage, and survival through multiple compression (the model pays 4.0x and exits at 3.5x); not “a few hundred basis points” (a few percentage points) of improvement. A seller should understand this model well enough to see which assumptions their evidence can strengthen, and which entry prices no operational program can rescue.

SECTION 08

Transaction Timing: Invest, Hold, or Transact

The decision should be based on the data, not the calendar

The stages of company growth are unchanged: validating, scaling, then maturing. Below roughly \$3M of ARR, a company is usually still validating, proving the product works and sells repeatably, and it is early to run a process. From \$3M to \$15M with strong growth, a company is scaling, and that is the natural zone for a full process, because enough buyers compete to move the price. From \$15M to \$30M, a maturing company meets buyers who weigh profit alongside revenue rather than revenue alone. AI has changed the timing decision inside each stage. The question is whether your evidence of durability is building faster than the competitive risk to it.

The Triggers

TRIGGER**Accelerate a process**

When the customers you lose or fail to expand are citing AI-native alternatives, when win rates are falling or you are cutting price to compete, and when closing the gap would take longer than the window is likely to stay open. Value is higher now than it will be once the erosion shows plainly in your numbers, and waiting does not improve it.

TRIGGER**Invest before a process**

When your moat is measurably widening, when customers are staying and spending more, when the roadmap that deepens your defensibility is funded, and when you can produce the evidence file (data rights, unit economics, priced expansion) within twelve to eighteen months. Selling now would sell the durability low, before it is provable.

TRIGGER**Hold and compound**

When your data position, your grip on customer workflows, retention, and unit economics are all improving faster than the competitive risk to them, and the expansion case described in Section 04 is genuinely yours to capture. Durable leaders in an expanding market should not be influenced by a market-wide narrative.

SECTION 09

AI Defensibility Case Study: Grow Faster vs. Grow Stronger

A classic tortoise-and-hare story: headline metrics favor the speedy hare, but the more durable tortoise wins the race

COMPOSITE A: "TORTOISE," VERTICAL COMPLIANCE SAAS

Tortoise has \$8M of ARR growing 18%. NRR is 108% and rising; GRR is 93%. Gross margin is 74%. The product is the system of record for regulated case files, wrapped in heavy integrations and audit-driven services. Its users are compliance officers making high-stakes filings, and usage is daily and organization-wide. Data rights are contractually secured, the product measurably improves with tenant data, and no AI-native competitor has surfaced in a renewal conversation.

COMPOSITE B: "HARE," HORIZONTAL WORKFLOW AUTOMATION

Hare has \$12M of ARR growing 32%, the better headline. But NRR is 112% and decelerating from 118%; GRR is 84% and slipping, with exit interviews citing AI-native alternatives on price. Gross margin is 68% with rising inference costs and no per-unit model. The services layer is thin, the product is not a system of record, its users are mid-level operators, and usage is broad but shallow. Two competitive AI mentions surfaced in the last renewal cycle.

Reading Tortoise. A strictly formulaic valuation prices Tortoise as a decent but unremarkable asset, with growth as the weakness. The AI defensibility analysis reaches the opposite conclusion. Tortoise is protected on all four dimensions: it is the system of record for regulated case files, heavy integrations and audit-driven services mean other systems depend on it, a mistake in a compliance filing is expensive, and its users are compliance officers doing high-stakes work rather than casual users. Its readiness is improving. And retention is getting stronger, with NRR at 108% and rising while GRR holds at 93%.

Every trigger in Section 08 points to investing before running a process. Twelve to eighteen months of work would go into converting the AI roadmap into priced expansion modules that lift NRR, and documenting the data moat so a buyer can verify it. That work addresses growth, the one weak spot, while the moat keeps compounding underneath it. Tortoise's realistic buyers are strategic consolidators and PE platforms, and both pay for what it is building: the system-of-record position, the contractual data rights, and the daily organization-wide usage. Transacting now would mean pricing Tortoise on evidence it has not yet produced.

SECTION 09

Reading Hare. The formulaic valuation prices Hare above Tortoise, on more scale and more growth. The AI defensibility analysis reaches the opposite conclusion. Hare is exposed on all four dimensions: it is not the system of record, few other tools depend on it, the cost of a mistake in its workflow is low, and its users are mid-level operators rather than people making high-stakes decisions. Its readiness to capture the AI shift is only partial. And the triggers are firing. Customers are leaving and citing AI-native alternatives on price, compute costs are rising without per-unit pricing to cover them, and NRR is decelerating from 118% to 112%.

Fixing this would mean rebuilding the product around a data loop that improves with customer use, and deepening the services layer so the product is harder to unplug. That work would take longer than the competitive window is likely to stay open. So Hare should accelerate a process now, while the growth story is intact and the erosion has not yet worked its way into the reported numbers. The buyers to run it toward are strategic consolidators and PE add-ons. They want the customer base and the workflow footprint, not standalone potential. This is not a fire-sale conclusion. Hare's best price arrives earlier than its owners would expect.

The takeaway: a static, formulaic valuation makes Hare look like the better business. It is measuring where each company stands, not where each one is heading. Tortoise is accumulating the evidence buyers pay for, and Hare is losing it. So Tortoise should invest before selling, and Hare should sell before the erosion is plain to everyone.

SECTION 10

The 85 Advisors Forward View

Our conclusions, and a suggestion or two

Metrics set the starting point; defensibility sets the confidence interval. A company with strong metrics and weakening defensibility is worth less than its metrics suggest. A company with softer metrics that is deepening its workflow depth and data position is worth more. Buyers are already pricing this, and an owner who does not know which side of that line the business sits on will be surprised by the feedback.

Work the two biggest levers, in the right order. Growth is the largest lever a company controls, and NRR is second, though NRR doubles as proof of defensibility. But processes are won on what the model equation cannot see: your defensibility evidence, the fit with a particular buyer, the quality of the process, and the structure of the deal.

Understand the buyer's model better than they expect you to. PE returns now depend on earnings growth after close, which buyers investigate hard and rarely pay for up front. Sellers who arrive with the operational evidence that makes that plan credible negotiate from better ground: on price where there is competitive tension, and on structure where there is not. The same burden applies to margins. If inference costs sit in your cost of goods, lower gross margins survive diligence only with proof that AI is central to your differentiation, that per-unit costs fall as you scale, and that the offsets are documented.

The bull case is real for some. Work out whether it applies to you. The research behind this report argues that fear of a "SaaS-pocalypse" badly understates what enterprises are gaining from AI, and that digital labor is a larger opportunity than any previous enterprise technology shift, with well-positioned incumbents among the winners. For a genuinely defensible company, the best path may be to hold and compound into that expansion.

The window that is open now will not stay open. In most sectors, clearly AI-defensible companies are not yet fully priced for their position, and vulnerable ones can still transact on their financial story. Both of those close as buyers move from treating AI as interesting to treating it as required. Know which company you are, then execute accordingly.

SECTION 11

What We're Watching, and What Could Change

Three things that would change our view, and what we would do about each

This report rests on judgments that could turn out to be wrong, so we track the evidence that would tell us. The three indicators below are the ones with the most leverage over its conclusions. Inference costs determine whether the margin standards we describe hold or loosen. The pricing of consumption revenue determines whether a company should move before a pricing pivot or through it. And whether buyers actually score defensibility in diligence determines whether our overlay describes how deals are priced today or how they may be priced later. Here is what we are watching in each, and what we would change if the evidence moves. Where it moves against us, we will say so in the next edition of this report.

INDICATOR	WHAT WE'RE WATCHING FOR	WHAT WE WOULD CHANGE
Inference-cost trajectory	Per-unit inference costs continuing to fall rapidly versus plateauing for production workloads	If costs keep falling faster than prices erode, we would ease the proof we ask for from companies carrying inference in their cost of goods, and allow exposed companies more time before thin margins become a real problem. If costs level off instead, we would raise the bar and expect a clear account of what each unit of usage costs to serve, and what happens to that cost as volume grows.
Hybrid-revenue pricing in processes	Whether companies with material consumption revenue clear processes at multiples comparable to subscription peers, or at a persistent haircut	If consumption-heavy companies clear processes at multiples close to their subscription peers, we would treat the increase in the proportion variable consumption revenue as one factor among many rather than an explicit M&A timing trigger (i.e., to sell while it is less of the revenue mix). If a lasting discount shows up instead, we would treat how a company earns its revenue as a valuation driver in its own right, weighed alongside growth and retention.
Defensibility scoring in live diligence	Whether structured AI-quality scoring shows up in actual buyer diligence requests, versus remaining a research construct	If buyers start asking these questions in diligence, we would fold these dimensions into standard preparation, the way we already do with retention and cohort data. If they have not by our next edition, or by mid-2027 at the latest, we will soften our stance that AI defensibility is priced into deals today and describe it as where pricing is likely heading.

— REFERENCE MATERIAL

Appendix

APPENDIX

AI Defensibility Diagnostic

Eight questions a buyer will ask, and what your answers should look like

The point of the exercise is to work through these questions now, while you have time to change the answers. There is no score. What matters is your ability to answer each one with the evidence a buyer will ask you to produce on each question. If Questions 2, 3, and 8 all come back weak, the timing conversation in Section 08 belongs on your near-term agenda.

QUESTION	STRONG LOOKS LIKE	WEAK LOOKS LIKE	EVIDENCE TO PRODUCE
1. System of record: if your product went down for two days, what stops in the customer's core business?	Work stops. Filings miss deadlines, transactions queue up, and someone escalates within the hour, because your product holds the authoritative record.	Customers keep the real record in a spreadsheet or another system, and yours is a place they copy things into. A two-day outage is merely an inconvenience.	Outage-impact narrative; integration map
2. What can a general-purpose AI assistant not replicate about your core function today?	Your answer points to an asset you own: proprietary data, deep integrations into customer systems, or a compliance position that took years to earn.	The answer is a list of features, and each one reads like something a competitor could ship next quarter, or that an assistant could approximate today.	Side-by-side task evaluation vs. a frontier model
3. NRR over the trailing four quarters, and its direction?	Above 110% and rising, with the expansion coming from customers buying more rather than from price increases.	Below 100% and falling. This is usually the first place displacement shows up in the numbers, months before it shows up in growth.	Customer-level NRR recalculation; cohort expansion detail
4. How much value sits outside pure software (services, compliance, integrations, data access)?	Leaving means re-implementing, retraining staff, and redoing compliance work. A customer would need a project plan and a budget to switch.	A customer could export their data and be running on something else within a week, with nobody outside the team noticing.	Switching-cost inventory; implementation timelines
5. Who uses the product, and what is at stake when they do?	Senior people use it for decisions or transactions that carry real money, and an executive protects the budget line personally.	Junior staff use it for low-stakes tasks, and it is the kind of tool that gets cut first when someone reviews the software spend.	User-seniority mix; contract-value distribution; decision-value examples
6. How embedded is the product in daily operating rhythm?	Used daily by several teams, and it is where the work actually happens rather than where it gets reported afterward.	Used weekly or monthly, mostly for reporting. If it went away, the work would continue and the reporting would move somewhere else.	Usage telemetry: frequency, breadth, depth by cohort
7. Is the product getting measurably better as customers use it?	You can show the product performing measurably better than it did a year ago, and tie that improvement to the growth of your installed base.	The product works the same for a customer in year three as it did on day one. Usage generates data, but nothing in the product improves because of it.	Model-performance deltas with and without proprietary data; data-rights file
8. Have AI-native competitors appeared in deals or renewals in the last twelve months?	None have come up, and customers are asking you to build more AI rather than asking whether someone else already has.	They are appearing in competitive deals, in requests for price concessions, or in the reasons given for losses. Any repeated pattern counts.	Win/loss log; renewal-conversation notes

Questions 1, 4, and 5 track SaaS Capital's SoR, NSC, and U&U dimensions respectively (Question 5 deliberately tests seniority and stakes, per the source).

Question 6 is 85 Advisors' workflow-embeddedness dimension.

DEFINITIONS

Definitions Used Throughout

Terms are calculated inconsistently across the market, best settled before a process

TERM	DEFINITION
AI COGS / compute	Inference and compute treated as a recurring cost of goods, tracked separately from ordinary infrastructure, with cost per unit of usage and its trajectory documented.
ARR	Annualized committed recurring revenue. It is becoming increasingly important to state explicitly whether consumption or usage revenue is included, and how that compares to the comps a buyer will use.
Customer concentration	Largest-customer and top-ten concentration measured separately, with ARR as the denominator. Present both; a hyperscaler relationship can carry strategic value and dependency risk at once.
Gross margin	Revenue less cost of goods sold. For AI-embedded products, separate inference/compute cost from implementation and human-delivered work, and evaluate contribution margin per unit of usage rather than the consolidated figure alone.
GRR / NRR	Gross and net revenue retention over trailing twelve months, computed from customer-level data: sum prior-period revenue for customers active in the prior period, then compare it to the same customers' current-period revenue. NRR includes expansion; GRR caps each customer at 100%. NRR is non-GAAP and inconsistently self-reported; recalculate it before benchmarking.
LTV:CAC	Gross-margin-adjusted lifetime value over fully loaded acquisition cost. Formulas vary materially, so present CAC payback months alongside the ratio, and note that a very high ratio can signal underinvestment in growth rather than operational excellence.
Recurring revenue %	The share of revenue that is genuinely recurring. State explicitly what you count as recurring versus consumption or usage-based, since hybrid revenue does not fit traditional ARR accounting and buyers will test the distinction.
Rule of 40	In this reference, ARR growth rate plus EBITDA margin. Market practice also uses free-cash-flow margin, and buyers will ask how capitalized development and adjustments are treated.
SaaS Capital 2026 Valuation Model Equation	Valuation Multiple = $-1.7 + (0.206 \times \text{SCI}) + (6.23 \times \text{ARR Growth Rate}) + (3.405 \times \text{NRR Rate})$; where SCI = SaaS Capital Index average revenue multiple. Its authors disclose that the model was estimated from 63 private valuation events, that its r^2 of 0.378 leaves most of the variation in observed multiples to factors outside the model, and that only about half of actual multiples fell within $\pm 30\%$ of its prediction.

SOURCES

Sources & Disclaimers

Research cited in this report

- SaaS Capital, "What's Your SaaS Company Worth?" (2026 update), published February 5, 2026: valuation formula, methodology (63 events, r^2 0.378, $\pm 30\%$ band), NRR calculation guidance, SCI usage cautions.
- SaaS Capital, "Introducing the SaaS Capital AI Assessment Framework," February 17, 2026: SoR / NSC / U&U definitions, rating methodology and Leader/Laggard clusters, stated scope exclusions.
- SaaS Capital Index: public B2B SaaS revenue-multiple index; cite the as-of date of any SCI value used.
- PitchBook, Q1 2026 Analyst Note, "Private Equity's Exposure to the Software Reckoning" (Garrett Hinds; published February 17, 2026): rate/multiple relationship and its breakdown, anchoring attribution, and the recent AI-attributed leg down ("software valuation reset" discussion); incumbent-defense and organizational-integration findings ("Enterprise switching costs" discussion).
- PitchBook, Q1 2026 Analyst Note, "SaaS Is Dead, Long Live SaaS" (published February 9, 2026): service-as-software thesis; software-and-payroll-budget convergence and TAM-expansion framing.
- PitchBook, Q1 2026 Analyst Note, "Through the Looking Glass: The Race to Build Enterprise AI" (published March 9, 2026): the eight evaluation frameworks for enterprise-AI readiness, including data-moat depth and agent-operations maturity ("eight frameworks" section and scoring methodology).
- PitchBook, Q2 2026 Analyst Note, "Mapping the AI Super-Cycle" (published April 24, 2026): service-as-software pricing transition, 2026–2031 (pp. 9–11); digital-labor expansion thesis and disagreement with the "SaaS-pocalypse" narrative (p. 7); incumbent transition risks (p. 39).

This report is prepared by 85 Advisors for general informational purposes. It is not a valuation opinion, a fairness opinion, or investment, legal, or tax advice, and no transaction outcome is promised or implied. It relies on third-party data believed reliable but not independently verified. Forward-looking statements involve uncertainty. Any valuation is specific to a company, a buyer, and a date. 85 Advisors is a sell-side advisory firm; readers should weigh recommendations accordingly.



Exemplary in our craft.
Exceptional in *your* outcomes.

85 Advisors
The Helmsley Building
230 Park Ave, 3rd Floor West
New York, NY 10169
United States

Tel: 646.854.3810 · Email: info@85advisors.co · www.85advisors.co