

— STRATEGIC REFERENCE

A Guide to B2B Software Company Valuations in the AI Era

*The ten quantitative and eight qualitative drivers
behind a software company's valuation — and how
AI shifts each*

INTRODUCTION

Our B2B Software Valuation Guide

Our revised valuation framework comprises ten quantitative and eight qualitative drivers, which we've 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 resource 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 overlaid 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.

How to Use This Reference

The pages that follow present the framework in three groups: quantitative drivers, qualitative drivers, and the new AI-specific diligence items. Each section includes a reference table followed by driver-by-driver commentary, with a closing section on the pricing-model transition now reshaping software revenue. In each table, **Weight** indicates how much the driver moves value, the **Tiers** are target levels, and **"What AI changes"** flags whether AI shifts the metric itself or only how a result should be read. The tiers and weights are directional judgment, not a valuation or a 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.

STRATEGIC REFERENCE

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. Customer 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.

STRATEGIC REFERENCE

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 are comparable to the comps a buyer will use, before a process begins. See the Pricing Transition section that follows for further discussion of the AI pricing-model transition.

3. ARR growth — amplified, and still the largest company-specific lever. Although the thresholds hold — 20% and above is strong, 30% and above is premium (see the tiers above) — 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, 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. From our observations, we can conclude that 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.

STRATEGIC REFERENCE

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 — service quality, customer health, pricing, implementation, segment mix, and macro conditions — remain. 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 months 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.

STRATEGIC REFERENCE

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.

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 later in this reference.

STRATEGIC REFERENCE

Qualitative Software Valuation Drivers — Commentary, Continued

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 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.

STRATEGIC REFERENCE

New AI-Specific Software Valuation Drivers

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

The following three topics did not exist as significant diligence items three years ago. They do now, and should be part of the checklist in preparing for any M&A process.

1. The Data Moat — And 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, collected at a point in time with no feedback from usage, is a head start that 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, not against a universal cutoff.

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.

STRATEGIC REFERENCE

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 exploit, or that customers can withdraw, fails diligence regardless of its quality. Buyers increasingly read the data-processing agreements before they read the model benchmarks.

A verified data moat changes which buyers compete and what thesis they underwrite. 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.

STRATEGIC REFERENCE

The Coming Software 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 introduces structural noise, and expansion and contraction both flow through it faster. We advise presenting cohort-level usage trends alongside the headline number to avoid overstating volatility.

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 underwrite hybrid revenue is evolving, and calibrating to current buyer practice is part of our standard pre-process work.

STRATEGIC REFERENCE

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.

This reference is provided for general informational purposes by 85 Advisors and does not constitute investment, legal, tax, or accounting advice, nor an offer or solicitation of any kind. Tier thresholds, weightings, and valuation benchmarks are directional judgments drawn from our experience and third-party research believed reliable but not independently verified; they are not a valuation or a prediction of any transaction outcome. Forward-looking views on AI and market conditions are subject to change. Consult 85 Advisors before acting on anything herein.



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