

# The Great SaaS *Reckoning*

How higher rates and agentic AI hit the software industry at the same time, and what Tier-1 LPs should do about it.

April 2026 | For LP use only

## EXECUTIVE SUMMARY

The software industry is getting kicked in the teeth from two directions at once, and the blows are landing at the exact same time.

On the financial side, the easy-money era that let sponsors load up software companies with debt at 15–20x revenue is officially over. Thoma Bravo just handed the keys to Medallia over to its lenders in April 2026, vaporizing roughly **\$5.1B** of equity. Vista took a similar bath on Pluralsight. When even JPMorgan can't get a **\$5.3B** debt deal for Qualtrics away because investors are openly muttering about "AI disruption," you know the old playbook is dead.

Meanwhile, agentic AI is quietly eating the seat-based pricing model that made SaaS so attractive in the first place. Klarna's AI assistant handled the work of roughly **700 humans** while managing 2.3 million conversations in its first month. If one agent can replace ten seats and you're still charging per seat — well, the math writes itself.

We're calling this the **Great Bifurcation**. On one side you have thick systems with real data gravity and regulated workflows that agents will likely augment rather than replace. On the other side are the thin, horizontal tools and shiny wrappers that are about to discover how fast "AI-powered" turns into "AI-replaced."

For LPs, the message is simple and uncomfortable: stop pretending every software asset is created equal. We need to be ruthless about tiering portfolios and applying real discounts to anything that still looks like classic seat-based SaaS with no credible path to outcome pricing. The Medallia wipeout may just be the preview.

# 1. Two Regime Shifts, One Industry

Between 2015 and 2025, private equity sponsors completed more than 1,900 software acquisitions, frequently at 15–20x revenue. The deals were predicated on three assumptions: near-zero borrowing costs, perpetual double-digit ARR growth, and seat count as a durable proxy for software value.

All three are gone at the same time.

Higher benchmark rates turned interest coverage ratios from comfortable to precarious across a large slice of sponsor-backed software. Where 5x levered buyouts required, say, 10% revenue growth to service debt, companies delivering 4–6% growth are underwater on coverage. Most of those companies are still pretending they aren't. The arrival of agents capable of automating entire workflows has meanwhile broken the seat-based pricing assumption that justified those multiples in the first place.

The combination is a structural reset (the **Great Bifurcation**) between thick systems that accumulate and act on proprietary data, and thin wrappers that are easily displaced. The market is already pricing the split. Most GP quarterly reports haven't caught up.

---

*"When JPMorgan and a syndicate of banks explicitly halt a \$5.3B debt sale because of 'AI disruption' and 'software pain,' that is a regime shift, not a blip."*

---

# 2. The Credit Reckoning

Lenders are no longer passive coupon-clippers in sponsor-backed software. They are restructuring agents, kingmakers, and in several recent cases, unwilling equity owners.

## MEDALLIA: \$5.1B GONE

Medallia is the headline everyone's talking about, and for good reason. Thoma Bravo took it private in 2021 at **\$6.4B**, financing roughly \$3B with debt at a time when leveraged-finance markets demanded almost no risk premium for software. By late 2025, the BDCs holding that paper were quietly marking it down: FS KKR Capital at roughly 79 cents, Apollo Debt Solutions at approximately 74. By April 2026, the pretense was over. Blackstone, KKR, Apollo, and Antares took the keys. Thoma Bravo's equity (and that of its co-investors) went to zero. Roughly **\$5.1B**, gone.

Was it purely rates? Not entirely. The business wasn't collapsing, but it also wasn't growing fast enough to service the debt pile once the cheap-financing window slammed shut. Add growing skepticism that legacy customer-experience platforms have much defensibility in an agentic world, and the refinancing simply wasn't there.

## PLURALSIGHT: THE J.CREW PLAYBOOK, REVERSED

Pluralsight told a similar story, just with extra drama. Vista Equity Partners bought it for \$3.5B in 2021. Pandemic-era growth proved the high watermark; as rates rose and corporate training budgets contracted, the topline went negative. Vista tried the classic J.Crew IP drop-down move in early 2024 to buy negotiating room with lenders. It didn't work. A Blue Owl-led group (including Ares, BlackRock, Goldman Sachs Asset Management, Oaktree, and Golub Capital) took control in August 2024, converting roughly \$1.2B in debt to equity. Vista and co-investors absorbed approximately **\$4B** in losses.


These weren't black swan events. They were predictable outcomes when you buy at peak multiples with peak leverage and then both things you were counting on (cheap debt and endless growth) disappear at the same time.

## QUALTRICS: WHEN THE DEBT MARKET SAYS NO TO AI RISK

In March 2026, JPMorgan and a syndicate of more than ten banks halted a **\$5.3B** debt financing for Qualtrics: a \$3.3B leveraged loan plus \$2B in high-yield bonds intended to fund the company's acquisition of Press Ganey Forsta. Leveraged-finance investors passed, citing anxiety about AI disruption in the customer-experience and employee-feedback software categories. Qualtrics' existing \$1.5B term loan, which had traded near par in February, fell to approximately 86 cents by mid-March. The signal is stark: credit markets are now explicitly underwriting AI displacement risk alongside leverage ratios.

## Equity *Wiped* in Landmark SaaS Buyouts

Approximate equity losses in USD billions — Medallia and Pluralsight restructurings, 2024–2026



**[Editor's Note]** Equity loss figures are sourced from Reuters, Bloomberg, and Private Equity Wire coverage of each restructuring. Medallia: ~\$5.1B loss on \$6.4B 2021 take-private (Thoma Bravo, April 2026). Pluralsight: ~\$4B loss attributed to Vista Equity Partners and co-investors on \$3.5B 2021 acquisition (Blue Owl-led consortium takeover, August 2024). Figures are approximate; final settlement terms may vary. All figures in USD billions.

### SYSTEMIC SIGNALS: BDCS AND THE \$46.9B DISTRESS PILE

Medallia and Pluralsight are the names that made headlines. The systemic picture is uglier.

As of late March 2026, BDCs were trading at their deepest discounts to NAV in more than five years. The median forward price-to-NAV stood at approximately **0.74x**, implying a 26% haircut. In some direct-lending portfolios, software accounts for 20–40% of sponsor-backed exposure. Bloomberg reported in early February 2026 that more than **\$17.7B** of US technology company loans dropped to distressed trading levels in a span of four weeks, the largest deterioration since 2022, bringing the total tech distressed debt pile to roughly \$46.9B, dominated by SaaS names.

UBS credit strategist Matthew Mish has projected \$75–120B in private credit and leveraged-loan defaults by year-end 2026. The numbers are messy and the range is wide, but the signal is clear: lenders have already figured out what a lot of GP quarterly letters haven't admitted yet.


### 3. The Valuation Reset

Public SaaS multiples have compressed sharply from the 2020–2021 peak. At that peak, growth-oriented SaaS indices traded at median EV/Revenue above 18x, with some concentrated indices exceeding 20x. By Q1 2026, the picture is different depending on which cohort you look at, but the direction is consistent across all of them.

| SEGMENT (Q1 2026)                                        | EV / REVENUE<br>MULTIPLE | NOTES                                                              |
|----------------------------------------------------------|--------------------------|--------------------------------------------------------------------|
| Public SaaS median (broad index, ~100 names)             | ~6.4x                    | Down from >18x at 2020–2021 peak                                   |
| Top-quartile public SaaS                                 | ~13.8x                   | High-growth, strong NRR, margin expansion                          |
| Bottom-quartile public SaaS                              | ~1.8x                    | Low growth, AI-exposed revenue, weak metrics                       |
| PitchBook Q1 2026 enterprise SaaS median (99 names, TTM) | ~3.3x                    | [Icarus Asia note] Lower universe; reflects more distressed cohort |
| Private lower-middle market (equity-backed SaaS)         | 3–6x ARR                 | Central tendency ~4–5x for decent-quality metrics                  |

## Public SaaS Multiple *Compression*

EV / Revenue multiples — 2021 peak vs Q1 2026 cohorts



**[Editor's Note]** 2021 peak figure (~18x) represents broad public SaaS median at the cycle top; some growth-only indices peaked above 20x. Q1 2026 figures sourced from multiple data providers including SaaSValuationMultiple.com, Windsor Drake, and PitchBook Q1 2026 Enterprise SaaS comp sheet. Data points reflect different underlying universes and should be compared directionally, not as precise equivalents. All figures are EV/Revenue (LTM revenue basis). [Icarus Asia estimate] for top and bottom quartile breakdowns based on aggregated public sources.

ICARUS ASIA

Sources: SaaSValuationMultiple.com, Windsor Drake, PitchBook Q1 2026

*Note: these are directional. Different indices give different medians, and the cohorts don't perfectly overlap. The compression trend is real regardless of which number you pick.*

The new variable in multiple determination is an implicit **AI defensibility filter**. Two companies with identical growth rates and NRR can trade at very different multiples depending on whether they look like thick Systems of Action with proprietary data moats, or thin wrappers at risk of agentic displacement. The market is bifurcating before the actual seat losses hit the income statement. Forward-looking investors are pricing disruption probability, not just current ARR.

WHAT THE MARKET ACTUALLY PAYS FOR TODAY

| METRIC                     | PREMIUM THRESHOLD               | TYPICAL MULTIPLE UPLIFT                                                                     |
|----------------------------|---------------------------------|---------------------------------------------------------------------------------------------|
| Rule of 40                 | <b>≥ 40%</b>                    | ~60–70% premium vs sub-40% peers                                                            |
| Net Revenue Retention      | <b>&gt; 110% (≥ 120% ideal)</b> | Multiple turns of ARR per incremental NRR point                                             |
| Gross Margin               | <b>75%+ (80%+ top quartile)</b> | Material at scale; compressed for AI infrastructure-heavy models                            |
| CAC Payback                | <b>&lt; 12 months</b>           | Signals sustainable growth economics                                                        |
| Revenue per Employee (RPE) | Steep upward slope              | Emerging AI-era signal; flat RPE with active agentic initiatives flags organizational bloat |

4. Agentic AI and the Empty Chair Crisis

The shift from passive LLMs to agentic AI systems is the structural shock to seat-based SaaS, and it's moving faster than most underwriting models assume. These agents don't just answer questions. They perceive, plan, act via tools, and adapt based on outcomes. They ingest data from CRMs, ERPs, and ticketing systems. They chain reasoning steps to break goals into subtasks. They call APIs, execute workflows, and coordinate with other agents. What actually matters: they write new proprietary data back to systems of record, creating feedback loops that compound their value over time. The seat becomes irrelevant once the agent has its own credentials.


This is why the distinction between **Systems of Record** and **Systems of Action** matters now in a way it didn't before. Systems of Record (CRMs, ERPs, HCM suites) are passive schema-driven stores waiting for humans to click buttons. Systems of Action sit on top of those records and orchestrate outcomes directly: resolving tickets, moving money, closing deals, reconciling ledgers. As the agent console replaces the SaaS dashboard as the primary user interface, the historical UI moat erodes.

THE KLARNA NUMBER

Klarna's AI customer service deployment, built on OpenAI, handled approximately **2.3 million customer conversations** in its first operational month. The company said it was doing the equivalent work of **700 full-time agents**: handling about 75% of chats, cutting resolution times from roughly 11 minutes to under 2 minutes, reducing repeat inquiries by approximately 25%, and projecting a roughly \$40M contribution to annual profit. Klarna retained human agents for complex escalations, which tempers the headline. But the unit-economics shock to seat-based vendors is not tempered: the revenue Klarna might otherwise have paid to a human-seat-priced CX software vendor did not rise with agent productivity. It should have fallen.

## The Empty Chair *Math*

Theoretical ARR impact on a seat-based vendor when agent displaces human seats



**[Editor's Note]** This is an illustrative model, not a confirmed forecast. It assumes a hypothetical vendor with 100 seats at \$10K/seat = \$1M ARR. Agent-displacement rates of 30%, 60%, and 90% are applied. "ARR at risk" assumes seat-based pricing with no compensation mechanism; "ARR retained" assumes a full pivot to outcome-based pricing that grows with agent volume. Real outcomes will vary by product category, contract structure, and customer negotiating leverage. [Icarus Asia estimate]

ICARUS ASIA

Illustrative model — Icarus Asia estimate based on publicly reported Klarna deployment data

## FROM SAAS TO SAS: SERVICE-AS-SOFTWARE

The coherent response for vendors facing seat erosion is a pricing pivot: from a per-user-per-month tax on human headcount to a per-outcome fee: per resolved ticket, per closed deal, per reconciled ledger entry, per prevented incident. We call this model **Service-as-Software (SaS)**. The software is no longer a tool a human uses; it is the automated service itself.

SaS pricing aligns vendor and customer incentives in a way seat pricing never did. The more work the agents do, the more revenue the vendor earns. Customers whose headcount shrinks still pay, and may pay more, as agent volume grows. Vendors successfully pivoting to SaS can grow revenue per customer even as human chair counts fall to zero.

We discount heavily any business where revenue is both predominantly seat-based and where credible agentic use cases exist but pricing remains unchanged. We reward businesses with early SaS pilots showing measurable economics: a defined percentage of contracts on outcome-based terms, demonstrable gross-margin uplift per automated workflow, and RPE inflecting upward as agents scale.

## 5. AI-Native vs Wrappers: A Practical Test

Every software company is "AI-native" now. Ask their sales reps. The term has been appended to so many pitch decks that it's essentially stopped meaning anything. We treat it as a specific technical and architectural claim that requires evidence, not a buzzword that gets a pass because it appears in the executive summary. The distinction between companies that actually own their AI layer and companies that are prompting someone else's model determines whether you have a durable moat or a product that disappears the next time Anthropic ships a new feature.

| CHARACTERISTIC      | AI WRAPPER                                         | AI-NATIVE SYSTEM OF ACTION                                           |
|---------------------|----------------------------------------------------|----------------------------------------------------------------------|
| Core architecture   | Thin UI or chatbot on someone else's model         | AI at the core of the execution engine, not a sidecar                |
| Proprietary data    | None or minimal; uses public model knowledge       | Owns a semantic layer mapping the business's entities and rules      |
| Feedback loops      | None or limited; does not write back to SoR        | Agents write back to systems of record, compounding proprietary data |
| Displacement risk   | Replaceable if model provider ships native feature | Switching costs rise over time as semantic layer deepens             |
| Valuation treatment | <b>-2-3x ARR penalty (Icarus Asia)</b>             | <b>Premium justified by data gravity</b>                             |

### THE FOUR-QUESTION WRAPPER TEST

For LPs and non-technical GPs evaluating a software holding, we propose four questions. Any "no" answer is a red flag:

1. If OpenAI or Anthropic shipped this feature natively tomorrow, would this company still have a defensible moat?
2. Does the company own its semantic layer and data model, or does it simply prompt a public model?
3. What percentage of value delivered depends on proprietary data, versus what the public internet already knows?
4. Do agents write back to systems of record and generate new proprietary data, or do they only read and display?

## 6. The Great Bifurcation: Thick vs Thin

The SaaS universe is splitting into two camps. Thick systems (ERPs, core banking platforms, industrial SaaS) have deep proprietary data, regulated workflows, and integration complexity that makes them expensive to replace. Thin systems (horizontal point tools, UI-centric apps, simple workflow wrappers) are exactly what agents can replace, automate around, or commoditize. The bifurcation is accelerating because agents are becoming capable of automating the workflows that thin systems were built to facilitate.

### OUR WORKING TIER FRAMEWORK (YES, IT'S BLUNT)

We've found that any GP who claims their entire software book sits in Tier 3 and 4 is either not looking hard enough or not being honest with LPs. Most portfolios have more Tier 1–2 exposure than the quarterly letters suggest.

#### TIER 1 — ALREADY DEAD

Equity is zero or close to it. Lenders own the company. These are credit situations now, not growth equity. We treat them as lessons, not investments, and as useful data points for what Tier 2 looks like in 18 months if nothing changes.

**See:** Medallia (April 2026), Pluralsight (August 2024)

*Icarus Asia stance: No new equity. Underwrite any recovery as credit. If you're still carrying these at cost, that's a conversation to have with your auditor, not us.*

#### TIER 2 — HIGH ALERT / PROCEED WITH EXTREME CAUTION

Still breathing, but highly levered, thin product, limited data moat, and no serious pricing pivot underway. Interest coverage is thin. These are the ones that keep us up at night. Most horizontal point solutions live here right now, whether their owners admit it or not.

**Analogues:** Horizontal point tools (scheduling, basic HR), legacy CX and employee-experience survey platforms

*Icarus Asia stance: We require a credible, time-bounded transition plan within 12–18 months: not a roadmap slide, an actual architecture. Model base-case exits at 2–4x ARR. Default is no new equity.*

### TIER 3 — IN TRANSITION

Real data gravity, agentic experiments actually running in production, and at least some contracts moving toward outcome-based pricing. Execution risk is still high, but there's a plausible path to survival. This is where the interesting work (and the follow-on capital) should be focused.

**Analogues:** Salesforce (Agentforce), ServiceNow, Workday with active agentic deployment

*Icarus Asia stance: Priority for follow-on capital when agentic traction is visible and measurable, not just announced. Track the % of workflows automated, the SaaS revenue mix, and whether RPE is actually moving. Model base-case exits at 4–8x ARR.*

### TIER 4 — THE REAL MOATS


Deep, regulated, messy data estates that agents will have a hard time replacing because the value is in the accumulated proprietary data, not the UI. Core banking, heavy industrial ERP, certain vertical SaaS with proprietary operational data. AI budgets flow toward these assets. These are the ones we actually want to own more of.

**Analogues:** Core banking platforms, heavy ERP (industrial, aerospace, regulated healthcare), industrial SaaS with proprietary sensor/operational data

*Icarus Asia stance: Core targets for new capital. Underwrite exits at 6–14x ARR depending on growth and NRR. Strategic buyers will pay premiums for data estates. Don't discount these just because the UI looks like it was built in 2008.*

## Illustrative NAV Exposure by *Tier*

Expected range in a generic software-heavy PE or late-stage portfolio — Q1 2026 [Icarus Asia estimate]



**[Editor's Note]** This is an illustrative model — [Icarus Asia estimate] — based on our assessment of a generic software-heavy PE or late-stage VC portfolio as of Q1 2026. Actual exposure will vary materially by manager, vintage year, and sector focus. "Any GP claiming no Tier 1–2 exposure is either under-disclosing or under-analyzing AI risk" is an editorial judgment by Icarus Asia, not a statement of fact about any specific fund. Ranges shown are midpoints; actual allocation bands are Tier 1: 0–10%, Tier 2: 10–20%, Tier 3: 35–50%, Tier 4: 30–50%.

## 7. How Winners Will Run: The Operating Playbook

### REVENUE PER EMPLOYEE AS A PRIMARY METRIC

Agentic organizations are targeting what some practitioners call "Department of One" models: replacing entire functional teams with small groups of human stewards orchestrating agents. Early data from AI-native marketing and CRM deployments shows roughly double the revenue per employee compared with traditional-staffed equivalents at the same ARR scale. Klarna-style contact centers, where a small lead team oversees what previously required hundreds of FTEs, are the logical endpoint of this trajectory.

We now treat **Revenue per Employee** as a central underwriting metric alongside Rule of 40 and NRR. We look for companies where RPE is clearly inflecting upward as agents scale, not just rising because hiring froze. Flat RPE with active agentic initiatives is a specific warning sign: automation gains are being absorbed by organizational bloat or misaligned incentive structures rather than flowing through to margin.

## **PRICING TRANSITIONS WE REQUIRE TO SEE**

We look for three milestones. First, a hybrid contract structure where a shrinking seat base is offset by rising outcome-based or usage-based components, meaning total ARR per customer holds or grows even as headcount falls. Second, pure outcome-based pricing in high-volume, lower-risk workflows where the output is measurable and unambiguous: tickets resolved, marketing touchpoints completed, transactions processed. Third, early evidence that customers are willing to pay more per outcome than they paid per seat: the key proof point that SaS is genuinely value-additive and not just a repackaging exercise.

## 8. LP Playbook: Diligence, Monitoring, Exit

This framework isn't just for Icarus Asia's portfolio. We use it on every GP in our allocation, and we'd encourage any LP to do the same. The questions below are ones we now bring to every AGM. Some GPs find them uncomfortable. That's usually a signal.

### QUESTIONS WE NOW ASK EVERY GP

- ❑ **Tier map, not company count:** Show us your software portfolio mapped across Tier 1–4 by enterprise value, not headcount of portfolio companies. A fund with 30 software names that are all Tier 2 and 3 looks very different from one with 5 Tier 4 anchors.

---

- ❑ **AI-native or AI wrapper — prove it:** For your top ten software holdings, which are true AI-native Systems of Action? What's the evidence? We want semantic layer ownership, agent write-back capability, proprietary data volume. "We have an AI roadmap" doesn't count.

---

- ❑ **The seat erosion problem:** What percentage of portfolio ARR is structurally exposed to headcount reduction by agents? Which companies have live outcome-based pilots, and what percentage of their ARR actually runs on those contracts today, not the projected figure?

---

- ❑ **Revenue per Employee trends:** Show us RPE over the last eight quarters. Is it inflecting upward as agents scale, or is it flat while the AI narrative runs ahead of the operational reality?

---

- ❑ **Refinancing risk:** Which portfolio companies would struggle to access debt markets without a materially different AI story? What happens if one or more BDC lenders mark to market and trigger covenant reviews in the next two quarters?

---

### MONITORING DASHBOARD — ICARUS ASIA QUARTERLY REPORTING

For our significant software positions, we will report quarterly across three dimensions:

**Financial:** ARR, growth, Rule of 40, NRR, gross margin, CAC payback, and Revenue per Employee trend.

**AI and Product:** Percentage of workflows agent-augmented versus fully orchestrated; seat-based versus outcome/usage-based revenue mix; percentage of value attributable to proprietary data and semantic layers versus generic model output.

**Risk and Credit:** Net leverage, interest coverage, covenant headroom where applicable, and qualitative AI disruption risk assessment by tier (1–4).

## EXIT FRAMEWORKS BY TIER

| TIER                           | EXIT MULTIPLE RANGE    | CONDITIONS                                                                                                                             |
|--------------------------------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <b>Tier 4 — Moat-Protected</b> | <b>8–14x ARR</b>       | Strategic buyers paying premiums for data estates and Systems of Action; top performers with strong AI-native operations               |
| <b>Tier 3 — Transitioning</b>  | <b>4–8x ARR</b>        | Requires visible SaaS pricing traction and agent-driven margin expansion before broad marketing; higher multiples as auditable upside  |
| <b>Tier 2 — High Alert</b>     | <b>2–4x ARR</b>        | Assumes lender influence and higher hung-process risk; focus on opportunistic sales, liability management, or selective restructurings |
| <b>Tier 1 — Wiped Out</b>      | <b>Credit recovery</b> | Treat as distressed debt situations; underwrite to recovery value, not growth multiples                                                |

## 9. Conclusion

Higher rates and agentic AI have ended the era of indiscriminate software investing. Levered, thin, seat-based SaaS built on the assumption that headcount and ARR would grow in lockstep forever. That model is over.

The opportunity runs in two directions. Defensively, use this framework to pressure-test every GP's software exposure before the next AGM. The Tier 1 losses are visible now. The Tier 2 losses are forming in portfolios where the AI transition plan is still a slide deck rather than a running architecture. Offensively, the same disruption that's destroying thin systems is creating a new category of winner: thick, data-rich systems with outcome-based pricing models that actually grow revenue as agents displace seats.

The Medallia and Pluralsight restructurings were predictable (in retrospect) by anyone willing to look at leverage, product thinness, and AI exposure at the same time. Most people weren't looking. We intend to.

# Appendix — Primary Source Verification

## KEY CLAIMS AND SOURCES

**Medallia restructuring (\$5.1B equity wipeout, April 2026):** Yahoo Finance / Investing.com; Bloomberg, April 2026; Reuters coverage of Thoma Bravo creditor handover.

**Pluralsight restructuring (~\$4B equity loss, August 2024):** Private Equity Wire; Bloomberg, August 2024; 9fin on Pluralsight J.Crew-style LME.

**Qualtrics \$5.3B debt deal halted (March 2026):** Bloomberg, March 17 2026; Seeking Alpha; Private Equity Wire.

**BDC discounts to NAV (~0.74x, Q1 2026):** Referenced in Bloomberg / Reuters coverage of Medallia restructuring and BDC earnings reports, Q1 2026.

**\$17.7B technology loans distressed (early February 2026):** Bloomberg Intelligence, February 2026. Total tech distressed pile: ~\$46.9B as reported. Note: source document attributes figure to January 2026; Bloomberg report covers a four-week window ending early February 2026.

**Klarna AI deployment (2.3M conversations, ~700 FTE equivalent):** Klarna press release and Klarna CEO public statements, February 2024. Resolution time and profit impact figures from Klarna-published data.

**SaaS multiples (Q1 2026):** SaaSValuationMultiple.com; Windsor Drake SaaS Valuation Multiples 2026; PitchBook Q1 2026 Enterprise SaaS Comp Sheet (median of 99 names at 3.3x TTM revenue as of March 31, 2026). Note: different indices produce different medians; 6.4x figure and quartile ranges are sourced from broader universe trackers. [Icarus Asia note: use directionally, not as single-source fact.]

**UBS default projection (\$75–120B in private credit by year-end 2026):** Matthew Mish, UBS head of credit strategy, as reported by Bloomberg and CNBC, February 2026.

---

## ICARUS ASIA RESEARCH · April 2026

This document is prepared by Icarus Asia for informational purposes and is intended solely for the use of the Tier-1 Limited Partners to whom it is addressed. It does not constitute investment advice, a solicitation, or an offer to buy or sell any security. All projections, estimates, and forward-looking statements involve material risks and uncertainties and should not be relied upon as predictions of actual outcomes. Past performance is not indicative of future results. Specific claims labeled [Icarus Asia estimate] or [Inference] reflect our analytical judgment and have not been independently verified by third parties. Recipients should conduct their own independent analysis before making any investment decision.