



C U B E   C A P I T A L

---

# **The Question Nobody Is Asking**

Why the Software Valuation Crisis May Be Worse  
Than the Dot Com Crash

---

Hani Iskander  
Partner, Cube Capital  
February 2026

## The Confession

Marc Benioff went on a podcast last September and said something that, in another era, would have been treated as a confession. Salesforce, the company that had built the modern software subscription model, had cut its customer support division from nine thousand people to roughly five thousand. Not because demand was falling. Because AI agents were doing the work. “I need less heads,” Benioff said. He was not apologetic. He sounded exhilarated. “It’s been eight of the most exciting months of my career.”

By February 2026, Salesforce had quietly cut hundreds more across marketing, product management, and data analytics. Workday cut 400 positions focused on customer operations. Oracle, Amazon, and Meta collectively eliminated tens of thousands of roles in the first months of the year. The pattern is unmistakable: software companies are using AI to reduce their own cost of producing and delivering software.

Now imagine you are sitting in a procurement office in Canberra or in the IT department of a large Australian bank. You have a Salesforce renewal coming up. You have just read that the company you are paying millions of dollars a year has publicly celebrated cutting half its support workforce, that its operating margins are expanding toward 36 per cent, and that its CEO considers the headcount reduction one of the highlights of his career. You do the simplest kind of arithmetic: if it costs them dramatically less to run this software, why does it cost us the same to use it?

That question, once asked, cannot be unasked. And it is being asked in procurement offices and boardrooms around the world right now. Zylo’s 2026 SaaS Management Index reports that organisations are aggressively auditing their software portfolios, increasingly demanding measurable return on investment before renewing contracts. L.E.K. Consulting and Constellation Research both forecast a significant shift toward outcome based and consumption based pricing in enterprise SaaS agreements over the next two years. The shift from “how many seats do you need?” to “what measurable value are you delivering?” is not theoretical. It is the conversation happening in the next renewal meeting.

But here is what most people have missed. The procurement question is not the real story. It is the first domino. And to understand what it sets in motion, you need to see a pattern that almost nobody is talking about, one that reaches far beyond Salesforce, far beyond Wall Street, and directly into the valuation of every software company in Australia.

If you are a founder or owner of a software company in Sydney or Melbourne with revenues between \$5 million and \$100 million, you might be tempted to view this as a problem for the big end of town. That would be a mistake. What happens to the valuation multiples of listed software companies does not stay on Wall Street. It cascades downward, through private equity benchmarks, through venture capital markdowns, through the comparables that acquirers and investors use to value private companies. And the cascade is not proportional. It is amplified. When Workday’s revenue multiple compresses by 40 per cent, your company’s implied valuation does not compress by 40 per cent. It compresses by more, because private companies carry an

illiquidity discount, because they lack the diversification and scale that partially insulate the large vendors, and because the forces I am about to describe strike the mid market with a directness that the global platforms can at least partially absorb.

In my previous paper, *The Third Great Reset*, I argued that the \$3 trillion software economy is being structurally rewritten. This paper asks the question that follows: what does that mean for the value of your company? The answer requires a story that most people have been telling themselves wrong.

## The Wrong Comparison

Everyone reaches for the dot com crash. The Nasdaq fell 78 per cent between March 2000 and October 2002. Over half of publicly listed internet companies failed entirely. Venture funding virtually disappeared for several years. And yet, within a decade, the survivors, Amazon, Google, Microsoft, became some of the most valuable businesses in the history of capitalism. The weak died. The strong recovered. Patience was rewarded.

It is a comforting parallel. It is also the wrong one.

In March 2000, Cisco Systems was the most valuable company in the world. It had real revenue, real profits, and a genuine monopoly on the networking equipment that powered the internet. When the crash arrived, Cisco's stock fell 89 per cent. But Cisco did not go bankrupt. It remained profitable throughout the downturn. It continued to dominate its market. By any operational measure, it was a well run company.

And yet, more than two decades later, Cisco's stock price has spent most of that time below its March 2000 peak.

The reason was not operational failure. It was something more insidious: structural commoditisation of its core market. Open source alternatives, cheaper competitors, and cloud architectures eroded the premium that had justified Cisco's extraordinary multiple. The company adapted, diversified, and remained profitable. But the market never again assigned it the valuation premium it once commanded, because the economics underneath had permanently changed.

That distinction, between a company that fails and a company that survives but never recovers its valuation, is the one that matters for understanding what is happening to software right now. And it explains why the dot com comparison is not just imprecise but actively misleading.

**In 2000, the problem was valuation without earnings. Companies were worth billions with no revenue. The crash was a correction of fiction. In 2026, the problem is the opposite: earnings without a future. The companies at the centre of the current selloff are real businesses with real revenue. Salesforce generates \$38 billion a year. Workday produces \$9.5 billion. Atlassian just posted its fastest quarterly growth in two years. They beat**

**earnings expectations quarter after quarter and still watch their share prices fall. The market is not punishing poor execution. It is repricing the business model itself.**

To understand why, and to understand what it means for your company, you need to follow the chain of consequences that begins with Benioff's confession.

## The Squeeze

Start with the obvious implication. When software companies use AI to reduce their own costs, their margins expand. Salesforce's operating margins have pushed above 20 per cent, with analysts projecting them reaching 36 per cent. In the short term, this is good news for shareholders.

But margins that expand because costs are falling create a problem that most vendors have not yet confronted. The savings become visible to their customers. Enterprise procurement teams are not naive. When a vendor publicly celebrates cutting half its support workforce through automation, the procurement department takes notice. When margins visibly expand while subscription prices remain unchanged, the procurement department does the arithmetic. And the arithmetic produces a question that has no comfortable answer: if it costs you dramatically less to produce this software, why does it cost us the same to use it?

The vendors are caught in what I think of as the margin trap. They cannot resist the cost savings AI offers without falling behind their competitors. But the savings they achieve become visible to their customers, who then demand that the savings be shared. Margins expand temporarily, then compress again as pricing catches up with costs. The net effect is not permanent margin improvement. It is revenue compression.

That alone would be significant. But it is not the real plot twist. The real plot twist is what happens next.

For decades, the enterprise software market has been stratified by customer size. Salesforce, Oracle, SAP, and Workday sold to large enterprises because the economics demanded it. The cost of configuring, implementing, and supporting an enterprise platform for a 50 person accounting firm or a 200 person logistics company was prohibitive relative to the revenue that customer would generate. So the mid market, roughly the \$5 million to \$50 million revenue segment, was left to smaller, more nimble software vendors who built simpler products at lower price points. Each inhabited a natural territory. The large vendors stayed upmarket. The smaller vendors owned the middle. It was a stable equilibrium.

AI dissolves that equilibrium.

When AI reduces the cost of configuring and deploying a platform to near zero, the economics of serving small customers change completely. A self configuring version of Salesforce or Workday becomes feasible for organisations that could never have afforded the traditional

implementation. The enterprise vendor can suddenly reach downmarket without the cost penalty that previously kept them out.

For the mid market software company, this is the threat that most founders have not yet connected. They see AI as a challenge to their product, which it may or may not be depending on their data moats and switching costs. What they have not considered is AI as a challenge to their *market position*. Not because their product becomes obsolete, but because a multi billion dollar incumbent with effectively unlimited resources can now serve their customers at a marginal cost that approaches zero. The competitive landscape above them shifts downward, and they are squeezed from a direction they were not watching.

At this point, a reasonable person might object. Embedded systems of record are genuinely sticky. Regulatory complexity protects incumbents. Multiyear contracts buffer the transition. Switching costs are real, and no amount of vibe coding eliminates overnight the deep integration that enterprise software builds into an organisation's operations. These are fair points, and they are the reason this is not a story of sudden collapse.

But Cisco had all of those advantages too. Cisco had real switching costs, real integration depth, and genuine technical superiority over its competitors. What Cisco could not defend against was the slow, structural erosion of the premium that its market position commanded. The revenue kept coming. The margins held. But the multiple compressed, permanently, because the market understood that the *future* growth rate had fundamentally changed. Survival and recovery are not the same thing. And that brings us to the force that I believe has the most profound implications of all.

## **The Product That Never Dies**

The traditional lifecycle of enterprise software follows a pattern so familiar that most people in the industry have stopped noticing it. An organisation identifies a need. It evaluates vendors. It selects a product. It then spends 12 to 18 months, and often millions of dollars, implementing that product: configuring it, customising it, integrating it with existing systems, training its staff. The product goes live. For the next five to ten years, the organisation lives with the product's capabilities and limitations, supplementing them with workarounds and customisations. Eventually, the gap between what the product does and what the organisation needs becomes intolerable, and the entire cycle begins again.

This replacement cycle is not a minor feature of the software economy. It is one of its primary growth engines. Every five to ten years, a significant portion of the market turns over. Organisations that once bought SAP replace it with Oracle. Oracle customers migrate to Salesforce. Salesforce users evaluate newer alternatives. The perpetual churn creates perpetual demand.

Now imagine a product that never falls behind.

A product that uses AI to continuously reshape its interfaces, workflows, and capabilities to match the evolving needs of each specific customer. Not a static artifact that accumulates limitations over time, but a living system that adapts. The customer's processes change, and the software changes with them. New regulatory requirements emerge, and the product reconfigures itself. The organisation grows, restructures, enters new markets, and the application follows.

This used to be science fiction. Now it is an engineering roadmap. Large language models can already interpret natural language descriptions of business requirements and generate functional application code. The step from "generate a new application" to "continuously modify an existing one" is not a conceptual leap. It is a problem being actively solved.

If software products become genuinely adaptive, something remarkable happens: the replacement cycle slows dramatically. An organisation that buys a truly adaptive accounting system, or manufacturing platform, or human resources suite may never need to buy another one. The product improves continuously. The implementation pain never recurs. The customer relationship becomes, in effect, permanent.

Which sounds like a vendor's dream, until you think about what it means for the industry.

Every financial model for a software company contains a number called total addressable market. That number is not static. It grows as new organisations form, as existing organisations expand, and, critically, as the replacement cycle generates recurring demand from organisations that already own software but will eventually replace it. The replacement cycle is the hidden engine of total addressable market growth. Without it, the addressable market for any category of software converges on the total number of organisations that need that category, multiplied by the price. Once every organisation that needs a CRM has one, and that CRM never needs replacing, the market is saturated. Growth reverts to the rate of new business formation and population growth, which in mature economies is low single digits at best.

**This is the structural risk that current valuations have not yet priced in.** The market has repriced software for AI agent displacement, for vibe coding, for the shift away from seat based pricing. It has not yet repriced for the possibility that the growth rate of the entire sector may have a structural ceiling far lower than anyone's current models assume.

Follow the full chain. The margin trap compresses revenue per existing customer as procurement teams demand that AI driven cost savings be shared. The downmarket invasion intensifies competition in the mid market as enterprise vendors offer simplified products at lower price points. Adaptive products extend customer lifetime but eliminate the replacement cycle that generates new demand. And at the end of the chain: an industry whose total addressable market stops expanding. Not because the products fail, but because they succeed too well.

## Five Companies on the Spectrum

Not every company is exposed equally. The five most prominent enterprise software companies each sit at a different point on what I think of as the vulnerability spectrum, and the differences are instructive for any founder trying to assess their own position.

### Atlassian

Atlassian sits at the most exposed end. Its core products, Jira and Confluence, are horizontal collaboration and project management tools. They have high adoption but relatively low switching costs, limited proprietary data moats, and functionality that is increasingly replicable by general purpose AI tools and vibe coding platforms. As of mid February 2026, Atlassian has fallen roughly 80 per cent from its 2021 peak despite posting record revenue. Both founders have been selling shares in consistent, visible blocks. The stock trades at a fraction of its historical revenue multiple. Analysts still rate it a buy with price targets that imply near doubling, but the gap between analyst consensus and market sentiment has rarely been wider.

### Workday

Workday occupies the middle ground, and it illustrates the tension between structural protection and structural vulnerability. Its human resources and financial management platform is more deeply embedded than Atlassian's tools, with genuine data moats around sensitive employee and financial information. Systems of record that hold payroll, compliance, and personnel data have real switching costs that no amount of vibe coding eliminates overnight. But Workday remains fundamentally seat based, and the CEO transition in February 2026, with co founder Aneel Bhusri returning to replace Carl Eschenbach, signals strategic urgency. Workday is pivoting toward consumption based pricing through what it calls Flex Credits, and has acquired the AI company Sana for approximately \$1.1 billion. The stock has fallen roughly 45 per cent in twelve months, trading near its 52 week low. The question for Workday is whether its system of record status provides enough structural protection to offset its pricing model vulnerability.

### Salesforce

Salesforce is the most complex case. It is simultaneously the poster child for the margin trap, having publicly reduced headcount from 9,000 to 5,000 in support alone, and a potential beneficiary of the agent economy through its Agentforce platform. The stock has fallen roughly 30 per cent from its highs, trading near \$190 at roughly 15 times earnings, down from 32 times a year ago. Salesforce's future depends on whether it can transition from selling software seats to selling AI agent orchestration, a pivot that would position it on the right side of the structural shift. Whether it can execute that transition while its existing customer base demands the pricing concessions that flow from its own cost reductions is the central strategic tension.

## Oracle

Oracle sits in a different position. Its database and infrastructure businesses give it exposure to the growth of AI workloads, not just the disruption of application software. Its cloud infrastructure division has been growing rapidly as enterprises deploy AI systems that require massive computing power. Oracle's application software faces the same structural pressures as its peers, but its infrastructure business provides a hedge that pure application vendors lack. The risk for Oracle is that it becomes a company with a thriving infrastructure division subsidising a slowly declining applications division, a duality that the market may struggle to value cleanly.

## Microsoft

Microsoft is the outlier. Its cloud infrastructure business, Azure, benefits directly from the AI computing boom. Its investment in OpenAI positions it at the centre of the AI platform layer. Its application software, the Office and Dynamics portfolios, faces the same structural pressures, but Microsoft is attempting to outrun those pressures by adding AI functionality through Copilot and charging a premium for it. The company recently raised Microsoft 365 prices, the second increase in four years, explicitly citing AI capabilities. Whether enterprise customers accept those increases or resist them will be one of the defining commercial negotiations of the next two years.

## Two Kinds of Survivors

The dot com crash produced a binary outcome. Companies without real businesses disappeared. Companies with real businesses eventually recovered and in many cases exceeded their previous valuations.

The current repricing will produce something more nuanced. All five of the companies I have described will survive. They have real revenue, real customers, and real cash flows. The question is not survival. It is whether their valuations will recover to previous levels. And the answer depends on which side of a structural divide they occupy.

On one side: companies that own unique data, that sit at the infrastructure or orchestration layer, that can transition their revenue models from seats to outcomes, and that can build the adaptive products that lock in permanent customer relationships. These companies will recover and may eventually command higher valuations than they do today.

On the other side: companies that are primarily application interfaces sitting atop commodity functionality, that remain dependent on seat based pricing, that lack proprietary data advantages, and whose products can be replicated by a sufficiently capable AI model. These companies will remain profitable. They will continue to generate cash. But they will become, in valuation terms, the next Cisco: operationally sound, strategically diminished, permanently repriced.

**Which side of this divide your company occupies is not determined by your current financial performance. It is determined by the structural characteristics of your business: your data, your embedding, your pricing model, your competitive moat.**

## **What This Means for Australian Software Companies**

For Australian technology companies in the \$5 million to \$100 million revenue range, the implications are acute. These companies are too small to compete at the infrastructure layer. They lack the resources to build the kind of AI orchestration platforms that might position larger vendors on the right side of the divide. Their natural market, the mid market segment, is precisely the territory that enterprise vendors will invade as AI reduces implementation costs.

The strategic options available to these companies are narrowing, and the window for exercising them is shorter than most founders appreciate.

Those with genuine structural moats, proprietary data, regulatory embedding, deep vertical expertise, high switching costs, should be deepening those moats now, not waiting to see how the market evolves. The founders who use this window to make their businesses indispensable will be positioned to command premium valuations in a consolidated market.

Those without structural moats face a more difficult calculation. The window for selling at the multiples of 2021 through 2024 has closed. But the window for achieving a strong outcome in the consolidation wave of 2026 through 2028 is open, and it will not remain open indefinitely. Once enterprise vendors begin reaching down into the mid market in earnest, the competitive position of unprotected companies deteriorates rapidly, and so does their value to potential acquirers.

I have met with founders who are only now beginning to confront these questions. For most, the greatest obstacle is not a lack of options but a lack of clarity about where they actually stand. The assumptions that guided their strategy for the past five or ten years were formed in a world that no longer exists. Testing those assumptions against the structural reality described in this paper is not optional. It is urgent.

## **The War That Matters**

Will software valuations recover?

Some will. The companies that sit at the infrastructure layer, that own the data, that orchestrate the agents, that build the adaptive products. For these companies, the current repricing will look, in retrospect, like an opportunity.

Most will not. Not because they fail, but because the growth assumptions that justified their valuations have been permanently altered by forces that are structural, not cyclical. The margin trap. The downmarket invasion. The end of the replacement cycle. The saturation horizon. Each

would moderate growth on its own. Together, they describe a new equilibrium for the software industry: lower growth rates, lower multiples, and a competitive landscape that bears no resemblance to the one in which most of today's valuations were formed.

Cisco survived the dot com crash. It won every operational battle in the years that followed. It remained profitable, it maintained market share, and it continued to innovate. By any conventional measure of business performance, Cisco succeeded. But it lost the war that mattered: the multiple. The market looked at Cisco's core business, understood that its growth trajectory had structurally changed, and never again valued it the way it once had. Cisco was a survivor that still lost.

The founders who will navigate the current transition successfully are not the ones waiting for the market to tell them what has changed. They are the ones who have already seen the structural shift, tested their assumptions against it, and begun making the difficult decisions early enough to act from a position of strength rather than desperation.

The question is not whether the market will recover. The question is whether your business is built for the market that is emerging. And answering that question honestly, with the kind of external perspective that internal confidence so often obscures, is where the real work begins.

---

## About the Author

*Hani Iskander is Founding Partner at Cube Capital (cube.capital), where he advises Australian technology companies on M&A strategy, capital structure, and organisational readiness for structural transition. He has founded technology companies in Australia and the United States, led the Asia Pacific region for a NASDAQ listed software company, and spent four decades operating at the intersection of technology and capital markets. He can be reached at hello@cube.capital*

## Sources

Dot com crash data, including the Nasdaq's 78 per cent decline, is drawn from market data reported by CNBC, the Library of Congress, and academic analysis from NYU Stern. The Cisco Systems precedent, including its 89 per cent decline and prolonged failure to sustainably regain its 2000 peak, is documented by Wikipedia's List of Companies Affected by the Dot Com Bubble and corroborated by multiple financial data sources.

Salesforce workforce reduction data is sourced from CNBC (September 2025), Business Insider (February 2026), and TheStreet (February 2026). Marc Benioff's comments regarding headcount reduction are from The Logan Bartlett Show podcast as reported by CNBC and NBC Bay Area. Workday layoffs and CEO transition data are from Investing.com and TIKR analyst reports (February 2026). Atlassian share price, insider selling, and analyst data are from TIKR, Public.com, and StocksGuide.

Enterprise software pricing evolution data is sourced from Zylo's 2026 SaaS Management Index, L.E.K. Consulting's 2026 SaaS pricing analysis, PYMNTS.com, Constellation Research's 2026 enterprise technology predictions, and BetterCloud's 2026 SaaS industry analysis. Software sector valuation data, including forward P/E ratios and price to sales compressions, is from Morgan Stanley, FinancialContent, and Bloomberg market data as of mid February 2026.

## Disclaimer

This paper is provided for informational purposes only and does not constitute financial, investment, or professional advice. The views expressed are those of the author and do not necessarily reflect the views of Cube Capital or its affiliates. Readers should seek independent professional advice before making any investment or strategic decisions.