



C U B E C A P I T A L

The Last Implementation

*Why AI-Adaptive Software Will End the Replacement Cycle
— and What Software Founders Must Decide Before It Does*

Hani Iskander

Founding Partner, Cube Capital

February 2026

© 2026 Cube Capital Pty Ltd. All rights reserved.

THE LAST RENEWAL

There is a saying that has circulated in technology circles for years, usually deployed as a warning against complacency: the Stone Age did not end because we ran out of stones. It ended because something better arrived, and the people still perfecting their flint tools never saw it coming. The enterprise software industry is standing in that moment right now. Not at the end of a product. At the end of a business model.

Consider the driverless car. For fifteen years, the automotive establishment, the insurance industry, urban planners, and transport economists treated full autonomy as an interesting aspiration, something to take seriously eventually, but not yet. The technology was always a few years away. The regulatory hurdles were always insurmountable.

Then Waymo began operating fully driverless commercial rides across designated service areas of San Francisco, with no human in the front seat and passengers in the back. Not a laboratory demonstration. A commercial service operating under approved conditions in a major city, in approved zones, at meaningful volume. The thing that everyone agreed was real in principle but distant in practice arrived ahead of schedule and without adequate warning for the industries it disrupted.

The taxi industry did not see the threat early enough to navigate it. Neither did the driving school industry. Neither, largely, did city planners. The pattern repeats: the disruption arrives before the incumbents have finished debating whether it is coming.

The enterprise software industry is in that debate now. And the disruption is closer than most vendors are publicly willing to say.

I want this to be wrong, for a simple reason: the last thirty years trained us to believe that great software, shipped faster, wins. That belief created fortunes. It also created a blind spot.

Replacement will still occur at the margins, for regulatory resets, vendor failures, and strategic platform shifts. What ends is replacement as the default growth engine in core categories.

Before you decide this sounds like the best version of SaaS, answer one question honestly: if the small group of customers that accounts for most of your ARR demanded a material price reset at renewal, would you still have the leverage to say no?

Here is the prediction I am prepared to make, and to be held to. Within five years, the dominant enterprise software products will not be periodically replaced by their customers.

They will evolve continuously, adapting to each organisation at a level of depth and specificity that makes the ten-year replacement cycle increasingly atypical in core categories.

I mean meaningful penetration, not universal dominance. Enough adoption to permanently reprice the growth assumptions on which current valuations depend.

The customer who implements a platform in 2026 will not look for a successor platform in 2031. They will simply continue using the one they have—by then so precisely calibrated to their organisation that it bears only a family resemblance to the product originally installed.

Now, if you are the chief executive or founder of an enterprise software company, your first instinct may be to find this prospect attractive. Customers who never leave. A permanent relationship. A recurring revenue stream that compounds indefinitely. It sounds like the best version of the SaaS model, taken to its logical conclusion. That instinct is exactly wrong, and understanding why it is wrong is the entire point of this paper.

When software adapts itself continuously, the SaaS licensing fee becomes very difficult to justify. If the product is improving itself using the customer's own data and processes, what exactly is the monthly rental paying for? When every implementation diverges from every other, the vendor's ability to standardise, to push one update to all customers simultaneously, and to own the resulting product IP unambiguously, comes under sustained pressure.

When the replacement cycle slows, the replacement-driven component of growth stops doing the work it once did. You can still grow through expansion and new use cases, but the easy reset-revenue of rip-and-replace becomes increasingly atypical.

And when only the largest companies can afford to build assurance-grade adaptive capability with defensible orchestration at enterprise scale, the economic conditions that made mid-tier software companies viable, and that made enterprise software the most reliable category in venture capital, quietly disappear.

Smaller vendors can still win, but only by choosing narrow adaptive wedges and deepening distribution or data advantages rather than attempting to match full-stack autonomy.

Put those pieces together and the shape is ugly: a revenue model that gets repriced, an IP dispute with no settled precedent, a market that saturates once replacement slows, and a tier of vendors that simply cannot fund the transition. In the categories that matter, those forces begin compounding within the next two renewal cycles—faster than most teams have planned for. Disruption rarely waits for consensus.

THE ACCUMULATION THESIS

Your customers already know more about how they use your software than you do. They just cannot articulate it yet.

Adaptive systems will articulate it, then propose changes, and increasingly implement them within defined guardrails, without a support ticket and without a roadmap debate.

Large language models can today interpret a natural language description of a business process and generate functional application code. They can analyse an organisation's existing workflows, identify inefficiencies, and propose modifications. They can read regulatory updates and assess their implications for a company's internal processes. None

of these capabilities, individually, is sufficient to produce a self-modifying enterprise application. Together, with integration into a production software environment, they describe exactly that.

The step from a static software application to an adaptive one has already been taken in the laboratory. What remains is an integration problem, and it is being actively solved by some of the best-resourced engineering teams in the world. Microsoft, Salesforce, SAP, and dozens of well-funded startups are all working on versions of what they variously call agentic AI, copilot functionality, or intelligent automation.

Each of these initiatives, in its most mature form, points toward the same outcome: software that does not merely respond to user commands but continuously observes how the organisation uses it, identifies the gaps between what the organisation needs and what the application currently does, and generates modifications that close those gaps.

Workflows. Integrations. Automations and interface conventions that exist nowhere else in the world. A platform that has become, in the most literal sense, theirs.

Think of what this means in practice. A law firm installs a legal practice management system. Over time, the system observes the way that firm handles matter intake, billing disputes, client communication, and regulatory filing.

It notices patterns the firm itself has never articulated: that certain partners consistently skip one workflow step, that client satisfaction correlates with a communication cadence that the system does not currently prompt, that a new court reporting requirement has created a bottleneck that a simple process modification would resolve. The adaptive system does not wait for the firm to raise a support ticket. It proposes the modifications. The firm accepts them. The application shifts.

This happens continuously, across every dimension of the application. The system never falls behind the organisation. It never accumulates the gap between capability and requirement that drives the replacement cycle. It simply adapts.

After several years of this process, something consequential happens: no two implementations of the same base product look remotely alike. The accounting system installed at a mining services company in Perth and the identical product installed at a logistics company in Melbourne diverge so far that comparing them is like comparing two humans who began from the same genetic blueprint. The shared ancestry is real.

The lived experience has made them entirely distinct.

The obvious response is that enterprise software has always been configured and customised, and that none of this is new. But adaptive systems do not just configure; they generate and accumulate organisation-specific logic at scale, blurring the line between implementation work and product IP in a way traditional customisation never did.

THE IP PARADOX

Here is where the prediction becomes genuinely uncharted territory, because it produces a legal and commercial paradox for which there is no settled commercial precedent at enterprise scale.

When Salesforce sells a CRM licence, the intellectual property is unambiguous. Salesforce owns the code. The customer licences the right to use it. This arrangement has been the legal foundation of the entire SaaS economy. It is written into every contract, reflected in every valuation, and assumed in every strategic plan.

Now imagine that same CRM after four years of continuous AI-driven adaptation. The system has generated thousands of workflow modifications, interface changes, automation rules, and integration adjustments specific to that customer. The base code is recognisably from Salesforce. But the accumulated modifications, the learned behaviours, the organisation-specific logic, who owns that?

The vendor will argue that the AI generating the modifications is their technology, therefore the outputs belong to them. The customer will argue that the modifications were generated by observing their proprietary business processes, using their proprietary data, to serve their proprietary operational needs. The modifications, on this view, are as much the customer's intellectual property as the vendor's.

The collision of two legitimate property claims is coming, and it will be litigated. More importantly, it will reshape the commercial relationship between vendor and customer before the first court judgment is handed down, because the customer who realises that their years of accumulated implementation investment constitutes a form of proprietary asset will begin demanding that their contracts reflect that reality.

This will not be decided by doctrine first. It will be decided by bargaining power at renewal.

Procurement will demand portability, escrow-like rights over customer-specific adaptations, and pricing that recognises co-created value. This quickly becomes concrete: ownership of workflow automations, portability of fine-tuned configurations and prompt libraries, and audit or export rights over learned behaviours.

The accumulated modifications, the learned behaviours, the organisation-specific logic after four years of adaptive AI: who owns that? This collision of two legitimate property claims will reshape software contracts before the first court judgment is handed down.

The implications reach beyond contract law. If each customer's implementation becomes a unique asset, the vendor loses something that has been central to their business model: the ability to standardise. Today, Workday can push an update to all its customers simultaneously. It can improve a feature once and deploy that improvement across its entire installed base. The economics of cloud software depend on this standardisation.

An AI-adaptive product, by definition, cannot be standardised in this way. The vendor is no longer maintaining one product. They are maintaining thousands of divergent variants of one product, each of which has evolved according to the specific needs of one customer.

The operational complexity of this is staggering. The commercial implications are worse.

THE PRICING COLLAPSE

Your renewal invoice arrives every year with the same justification buried inside it: we kept improving the product, so you keep paying for it. That logic has held for thirty years.

Adaptive software breaks that bargain from both ends.

On the vendor side, the justification for ongoing fees shifts from features shipped to assurance provided. The release cycle, the version update, the new feature announcement—these become background noise when the software is evolving continuously from within.

What remains is security, governance, auditability, uptime, and indemnity. Those have genuine value, but they do not command the margins of a product monopoly. Pricing follows the justification: from subscription to utility.

At the same time, the customer's embedded investment in the adaptive product becomes an argument against the entire rental model. If the system has spent four years learning your organisation, if it now reflects your specific workflows, your regulatory environment, your staff preferences, your client communication patterns, then the product is, in a meaningful sense, yours.

Paying a monthly rental fee for something that is effectively your own IP is a commercial arrangement that no sophisticated procurement team will accept indefinitely. Once AI drives measurable savings, procurement treats that saving as shared value, not vendor monopoly.

The pressure will not move toward elimination of fees. The vendor still provides infrastructure, maintains the underlying AI models, manages security, and provides the technical foundation on which the adaptation runs. Those services have genuine cost and genuine value.

But the pricing model will be forced toward something that looks less like a subscription and more like a utility fee or a maintenance contract: a flat or declining charge for infrastructure and base services, with the extraordinary margins of the SaaS era becoming an artefact of a model that no longer reflects economic reality. Value moves from licences to outcomes, assurance, and orchestration.

For a company like Salesforce, whose own guidance projects full-year revenue of approximately forty-one billion dollars at non-GAAP operating margins of thirty-four per cent, the implications of even a moderate repricing in this direction are severe. For a mid-tier vendor generating forty million dollars with margins that fund a development team and leave limited capital reserve, they are existential.

THE EXTINCTION GRADIENT

Not all software companies are equally exposed to these forces. The ones most likely to survive are the largest. The ones most likely to disappear are the mid-tier. That is almost precisely the inverse of what most people would assume, and understanding why requires looking at the forces clearly.

Microsoft, Oracle, SAP, Salesforce, and Workday have something that no smaller competitor possesses: the capital reserve to fund the AI infrastructure required to build genuinely adaptive products. The investment required to develop, train, and continuously improve the large language models and orchestration systems that will underpin adaptive software is measured in the billions.

A company generating twenty or fifty million dollars in revenue cannot get close to this threshold. Even a company at five hundred million dollars faces existential risk trying to cross it.

The medium and small application software vendors face a trap with no exit. They lack the capital to build the AI layer that would make their products adaptive. Building AI features on hosted models is cheap; building assurance-grade autonomous change inside regulated production systems is not. The distance between those two things is precisely where the capital threshold sits: on one side, an AI assistant layered on top of a product; on the other, a system that autonomously modifies production business logic within an auditable governance framework. Without adaptive capability of the second kind, their products will age and accumulate the gaps that drive replacement.

But the replacement cycle they depend on for growth may be ending precisely because the large vendors, who can fund the transition, are building the adaptive capability that mid-tier and small vendors cannot.

Which of those categories does your product actually sit in?

This is a consolidation dynamic unlike any the software industry has seen. Previous waves of consolidation were driven by M&A activity, by acqui-hiring, by the straightforward logic of larger companies absorbing smaller ones for their customer bases. What is coming is more fundamental: the structural elimination of the economic conditions that made small and mid-tier application software viable as a business category.

The structural elimination of the economic conditions that made small and mid-tier application software viable is not a wave of M&A. It is the logical consequence of an investment threshold that only the largest companies can cross.

The obvious objection to this argument is that switching costs have always protected incumbents, and that enterprise software is stickier than it appears. A hospital that has spent three years embedding a clinical management system into its operations, trained four hundred staff on its workflows, and integrated it with eighteen other systems does not

replace that software lightly, regardless of what the large vendors are building. The objection is fair, and the stickiness is real.

But it misunderstands the nature of the threat. The argument is not that existing customers will leave their current vendors abruptly. It is that the conditions which historically guaranteed a replacement purchase five to ten years later will no longer materialise. The hospital that does not leave will also not replace. It will simply continue, indefinitely, with a system that adapts to avoid giving it a reason to look elsewhere.

The revenue does not disappear overnight. The growth assumption disappears. And in software, where valuations are built almost entirely on future growth, the growth assumption is everything.

The venture capital ecosystem built around enterprise software startups faces the same logic from a different angle. Enterprise software has been one of the most reliable expressions of that thesis for three decades. A small team with a genuinely better product for a vertical market could capture meaningful share of a multi-billion dollar category and compound that position over years.

That thesis depends on three conditions: that there is a large market for the product, that a well-funded startup can compete meaningfully within it, and that the product can be maintained and improved with startup-scale resources. Adaptive AI disrupts all three conditions. The replacement-driven component of category growth converges toward saturation as the replacement cycle slows.

The competitive position of a startup against an incumbent whose product has four years of accumulated adaptation to each customer is almost impossible to establish. And the investment required to build a competitive AI layer is, by definition, beyond startup scale.

None of this means that software startups disappear entirely. Infrastructure plays, AI model development, specialised tooling for the orchestration layer: these represent genuine venture opportunities. But the category that built the most reliable venture returns of the past three decades, the vertical SaaS application serving a specific industry need, faces a structural threat that the venture community has not yet fully incorporated into its investment theses.

WHAT THE GIANTS MUST DO

If you run a large enterprise software vendor, the prediction I have outlined presents a peculiar strategic problem. The forces I have described erode the replacement-driven component of growth by weakening the replacement cycle, and compress pricing power by undermining the traditional SaaS fee logic.

And they will simultaneously require you to make the largest capital investment in your company's history to build the adaptive AI layer that might preserve your relevance.

The reason this is happening now, rather than in a decade, comes down to timing. Enterprise procurement teams have developed the sophistication to demand outcome-based value rather than feature counts. Security and compliance frameworks have matured enough to allow AI agents to operate inside production systems (something that was

genuinely not true three years ago). Model orchestration has caught up with the ambition. And the largest vendors now have distribution infrastructure that lets them reach downmarket at near-zero marginal cost. None of those conditions existed simultaneously before. Now they all do.

The instinct for most large organisations in this position is to defend the existing model while making incremental investments in the new one. This is almost certainly insufficient.

The companies that survive this transition will be the ones that cannibalise their own existing revenue models before someone else does: transitioning from seat-based to outcome-based pricing, investing in genuine adaptive capability rather than cosmetic AI features, and renegotiating intellectual property frameworks with customers to acknowledge the co-created value that accumulates over time.

Accepting the short-term pain of lower recurring fees in exchange for the structural position of owning the adaptive layer is not a comfortable decision. It is, however, the only one with a credible endpoint.

The broader pattern matters more than any individual vendor trajectory: the companies that survive will be those that can credibly charge for outcomes rather than access, and that own the assurance and orchestration layer on which adaptive software depends.

For Australian founders in the five to one hundred million dollar range, this pattern is not an abstraction. It is the competitive environment in which their next renewal, their next valuation conversation, and their next strategic decision will take place.

THE AUSTRALIAN QUESTION

In transactions and renewal negotiations, this is already changing behaviour. More renewals are starting to include AI cost-sharing clauses and demands for measurable outcome evidence. More acquirers are discounting valuations for products without demonstrable data moats or defensible workflow embedding. More founders are discovering that the multiples they assumed would apply to their next fundraise or exit have shifted beneath them.

The first casualty is not revenue. It is leverage: at renewal, in procurement, and in the valuation model your next buyer applies.

For Australian technology companies in the five million to one hundred million dollar revenue range, the analysis I have presented resolves into a single urgent question: which side of the investment threshold are you on, and what are you doing about it?

Most Australian application software companies are not on the right side. They do not have the capital reserve to build a competitive AI layer. Their products serve markets that the global giants, once they achieve adaptive capability, will be able to reach downmarket at marginal cost. And the window during which their specific expertise and customer relationships constitute a durable competitive advantage is narrowing faster than most founders appreciate.

The strategic options available to these companies are real but require honest assessment. Companies with genuine data moats, deep regulatory embedding, or specialised vertical knowledge that cannot be replicated by a generalist AI model have a defensible position, but they must deepen that position now, before the capability gap between their products and the adaptive giants becomes visible to their customers.

Companies without those structural advantages are facing a consolidation wave, and their optimal response is not to wait for a better offer but to engage the market while they still have meaningful leverage.

The response most commonly heard when these dynamics are raised is: yes, but this will take longer than you think. Perhaps. I have heard some version of this in most of the conversations I have had with founders over the past eighteen months, and I understand the instinct, because it is uncomfortable to plan around a disruption that has not yet fully arrived.

But the AI investment cycle does not wait for the market to be ready. The adaptive products being built today by the large vendors will begin reaching customers within two years. The contract renewals happening in 2027 and 2028 will be negotiated in an environment where the possibility of truly adaptive software is no longer theoretical.

And the valuation conversations happening between Australian software founders and acquirers in 2026 will reflect assumptions about the future that are being revised, quarter by quarter, faster than most people expected.

The question, then, is not whether this is coming. It is which of four positions your company actually occupies, because the right response is entirely different depending on the answer.

The first position is the one every founder wants to believe they occupy: a genuine structural moat. Proprietary data accumulated over years of customer use that no competitor can replicate. Deep regulatory embedding in an Australian context, where ASIC obligations, clinical governance requirements, or state procurement frameworks create switching costs that are genuinely prohibitive. Vertical expertise so specific that a generalist AI model cannot credibly serve the category without years of domain investment.

If this describes your business honestly rather than aspirationally, the task is not to panic but to fortify. Deepen the data advantage. Extend the regulatory integration. Make the switching costs more visible in your contracts. The founders in this position who act now, while the gap between their embedded value and the adaptive giants' capability is still wide, will find that the consolidation wave of the next two years rewards them rather than threatens them.

The second position is less comfortable but more common: strong distribution, loyal customers, and a product that customers genuinely like, but without a moat that would survive a determined assault from a vendor with adaptive capability and unlimited reach. These companies have leverage today that they will not have in thirty-six months.

Their customers are not leaving, but the next renewal conversation will be different from the last one, and the one after that will be different again.

For founders in this position, the most sophisticated decision available is to engage the acquisition market now, on their own terms, rather than later under pressure. A company with strong retention, a clean cap table, and genuine vertical relationships is attractive to a

larger domestic player or an international vendor seeking Australian distribution. That attractiveness is a depleting asset in the current environment, not a compounding one.

The third position is occupied by companies with genuine technical depth but limited market reach: a product that is genuinely hard to replicate but whose customers are concentrated, whose distribution is thin, and whose revenue has plateaued below the threshold that attracts serious acquisition interest independently.

For these founders, the most valuable near-term move is a channel or co-distribution arrangement with a larger platform, one that extends their reach before the adaptive giants fill the space their product currently occupies. This is a partnership conversation, not an exit conversation, but it needs to happen before the platform in question decides it no longer needs the partnership.

The fourth position is the one nobody wants to name aloud. A product that is good but not irreplaceable, in a market that is well-served, with competitors who are larger and better-capitalised. The honest counsel for a founder in this position is that the window for a strong outcome is open now and will not remain open indefinitely. Waiting for a better offer is a strategy that historically produces worse offers.

The founders who exit well from this position do so by accepting the reality of their structural situation early enough to act from strength rather than necessity.

Whichever of those positions describes your company, three decisions belong on your desk before the end of 2026. The first is a contract posture decision: what language governs the ownership of AI-generated adaptations and customer-specific model outputs in your existing agreements, and what position will you take in renewal negotiations?

The second is a product posture decision: what one capability, built on hosted AI models rather than proprietary infrastructure, would most meaningfully increase your switching costs and your customers' dependence on your accumulated knowledge of their operations? The third is a capital posture decision, and it is the most important: are you building to survive independently, or to achieve a premium transaction?

The answer shapes every other choice, and founders who straddle it do not get the benefit of either strategy. They get the costs of both.

The founders who navigate this transition well are not the ones with the most confidence in their products. They are the ones with the clearest view of their structural position, who can separate the question of whether their software is good from the question of whether their business model is durable in the world described above. Those are not the same question. Is the distance between them where your company currently stands?

The last implementation, the final time an organisation buys an enterprise application and then, five years later, replaces it with something else, may already have happened somewhere in the world. We will not know it was the last until several years from now, when the pattern is visible in retrospect.

But the strategic decisions that determine whether your company is on the right or wrong side of what follows are decisions that need to be made now, not in retrospect.

The question is no longer whether your product is good. It is whether your renewal leverage is compounding or depleting.

The founders who navigate this transition well will not be the ones who waited for the market to confirm what the structure already shows. They will be the ones who sought an honest external view of their position early enough to act from strength rather than necessity. That mandate, typically at board or shareholder level, when the decision is no longer theoretical, is the work Cube Capital is retained to do.

The craftsmen who were still perfecting their flint tools when bronze arrived were not failing at their work. They were simply operating inside a model that had already been superseded. They were the last to know it.

About the Author

Hani Iskander is Founding Partner at Cube Capital (cube.capital), where he is retained by founders and boards of Australian technology companies on confidential mandates involving 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.

Contact: hello@cube.capital