

# 'SaaSpocalypse': A Possible Framework for Assessing AI Disruption Risk

How do structural moats determine which software businesses are most vulnerable to the AI-native challenger wave

The recent wisdom in technology investing is that artificial intelligence threatens the entire software stack. In this telling, any vertical SaaS product is only a few foundation model iterations away from obsolescence, replaced by an AI-native competitor that can replicate its core functionality at a fraction of the cost and with none of the legacy technical debt. This narrative is not without merit.

AI is already enabling new entrants to compress years of product development into months, and to undercut incumbents on price in ways that would have been structurally impossible three years ago.

But the narrative is dangerously blunt. It treats all SaaS revenue as equally fragile, all customers as equally price-sensitive, and all switching decisions as equally unconstrained, none of which are true. The reality is that the risk of AI-driven disruption varies enormously across the software landscape, and the factors that determine vulnerability are knowable, measurable, and investable.

This analysis sets out a framework for differentiating AI disruption risk across the SaaS and general software universe.

We identify ten structural characteristics that meaningfully insulate an incumbent SaaS provider from the AI-native challenger wave and explain why each one matters.

## The Disruption Threat in Plain Terms

Before assessing resistance, it is worth being precise about the threat.

AI-native SaaS competitors enjoy several structural advantages over incumbents.

1. They can build core workflow automation at materially lower cost, leveraging foundation models rather than bespoke engineering.
2. They face no internal bureaucracy protecting legacy architecture.
3. They can price aggressively because their cost to serve is lower, and

4. They typically offer more intuitive interfaces built around natural language interaction rather than the complex menu-driven UIs that characterised the previous generation of enterprise software.

The categories most exposed are those where the core product is essentially a database plus a workflow engine with a user interface — and where

- the customer relationship is shallow,
- the data is generic, and
- switching costs are low

Think basic HR platforms, generic project management tools, or simple CRM systems where the customer data is portable and the product's value is largely replicable.

But much of the SaaS universe does not look like this. For a significant portion of enterprise software, the ten factors below create structural resistance to disruption that is durable, not transient.

## Ten Structural Moats Against AI Disruption

### 1. Legacy Industry Customer Bases

Customers in legacy industries such as construction, agriculture, mining, utilities and traditional manufacturing are structurally slower to adopt new technology regardless of its merits. Their procurement cycles are long, their IT teams are small and cautious, and their leadership is typically composed of domain experts rather than technology adopters.

For these customers, the question is never simply whether a new AI-native product works better; it is whether the organisation has the internal capability and appetite to evaluate, procure, implement, and support it.

An AI-native challenger entering these verticals faces not just a product competition but a sales and change management challenge. Incumbents with established relationships, trained customer support teams, and many years of implementation experience hold significant advantages that do not disappear because a competitor appears to have a better language model underneath.

## 2. Deep Vertical Integration and Specialist Domain Expertise

SaaS businesses that are deeply vertically integrated into a specific industry's workflows. Managing the full operational cycle rather than a single function are substantially harder to displace.

The value of these platforms is not simply their features, it is the accumulated codification of industry-specific logic, testing and dealing with edge cases, complying with regulatory requirements and security standards, and workflow nuances that took years of close collaboration with customers to build.

A new AI entrant can rapidly replicate the visible surface of a product. It cannot rapidly replicate the decades of specialist industry knowledge embedded in its data models, pricing logic, and exception handling.

The deeper and more specialised the vertical integration, the longer the runway before an AI challenger can match incumbent functionality credibly.

## 3. High Regulatory Burden and Compliance Complexity

Industries subject to heavy regulatory oversight such as financial services, healthcare, pharmaceuticals, legal, and energy, impose compliance obligations on their software vendors that represent a formidable barrier to new entrants.

Achieving regulatory certification or audit readiness is not a technical problem; it is an organisational, legal, and operational challenge that takes years and significant capital to address.

For the incumbent, this compliance infrastructure is already built. For the AI-native challenger, it represents a long lead time and substantial ongoing cost before it can even enter meaningful conversations with regulated customers.

This dynamic is particularly pronounced in jurisdictions with active enforcement, where procurement decisions require extensive due diligence on vendor compliance with such security standards.

## 4. Monopolistic or Oligopolistic Industry Structures

Customers that operate in concentrated market structures, such as regulated utilities, dominant national telcos, state-owned enterprises, large infrastructure operators, face fundamentally different incentive structures than companies in competitive markets.

Such customers are not fighting for survival on margin. The pressure to adopt lower-cost software solutions that characterises competitive industries is simply absent.

For these customers, the switching cost analysis is not about cost savings; it is about operational risk. An incumbent SaaS provider serving a large, regulated utility benefits from a customer that is genuinely indifferent to whether a competitor might offer the same capability at 30% lower cost. The utility is not optimising for software spend. The disruption risk of switching is a more salient consideration than the cost savings opportunity.

## 5. High Trust and Data Security Requirements

In sectors where the consequences of a data breach or system failure are catastrophic (think defence, intelligence, financial services, healthcare, critical infrastructure), trust is not a soft metric. It is a procurement prerequisite.

Customers in these sectors conduct extensive vendor security assessments, require specific certifications, maintain approved vendor lists, and in some cases are legally prohibited from using software that has not passed government security accreditation.

An AI-native challenger, regardless of its technical merits, cannot shortcut this accreditation process.

The incumbent, already on the approved vendor list with a track record of secure operations, holds a structural advantage that may take the challenger many years to replicate, during which time the incumbent can continue to evolve its product and further enhance its competitive moat.

## 6. Deep Integration Into Existing Technology Stacks

SaaS applications that sit at the centre of a customer's technology architecture; deeply integrated with ERP systems, data warehouses, operational databases and adjacent applications, create switching costs that are technical as well as organisational.

The cost of replacing such a system is not simply the licensing cost differential. It includes the cost of re-integrating with all upstream and downstream systems, retraining users, migrating data, and managing the transition period.

For enterprise customers with complex technology needs, this integration switching cost is often an order of magnitude larger than the annual software cost.

AI-native challengers that can demonstrate feature parity still face the uncomfortable reality that replacing a deeply integrated incumbent requires the customer to fund a multi-year transformation programme. This is a commitment most customers only make if the incumbent has materially failed.

## 7. Mission-Critical Application Status

When a SaaS application is genuinely mission-critical (meaning the business cannot function if it goes down), the risk calculus around switching changes fundamentally.

For these applications, the evaluation criterion is not 'is the new product better?' but 'can we be certain the new product will work as well as the one we already depend on?'. The asymmetry of this question strongly favours incumbents.

Clinical management systems in hospitals, trading platforms in financial services, dispatch systems in logistics, and operational control systems in utilities all fall into this category.

The consequence of a failed migration is not a productivity dip; it is a business crisis. This reality makes customer organisations deeply conservative about switching, regardless of how compelling the AI-native alternative appears in a demonstration environment.

## 8. Network Effects and Pooled Industry Data Benefits

SaaS businesses that aggregate non-public data across their customer base (such as industry benchmarks, cross-customer workflow analytics, anonymised operational data pools etc.) provide a form of value that is structurally unavailable to a new entrant by definition.

The insight a SaaS customer gains from understanding how their performance compares to anonymised industry peers, or from accessing benchmarks derived from thousands of similar organisations, requires a large installed base to produce. It's a chicken-and-egg quandary for the AI interloper.

This creates a compounding advantage. The more customers the incumbent has, the richer the network data, and the more valuable the platform becomes. This makes incumbent displacement more costly for each individual customer.

An AI-native challenger starting from zero cannot replicate this data pool regardless of its model sophistication.

## 9. Proprietary Non-Public Training Data

The quality of an AI system is determined in large part by the quality and relevance of its training data.

In many verticals, the most important data is not publicly available. It exists only within the systems of companies and institutions that have accumulated it over decades of operations.

Incumbents who have been processing, storing, and analysing this data for years are positioned to fine-tune AI capabilities on datasets that simply do not exist in the public domain.

This advantage is self-reinforcing. The more operational history the incumbent accumulates, the better its AI-augmented product becomes.

The new entrant, training on generic or synthetic data, will produce a product that may appear to perform well in general cases. However, a deeper examination of its capabilities will highlight its lack of domain-specific precision that enterprises in specialist fields quickly learn to notice and distrust.

## 10. Public Sector and Not-for-Profit Customer Bases with Asymmetric Procurement Risk

Perhaps the most underappreciated moat is the nature of the procurement decision-maker in public sector and not-for-profit environments.

Local government, higher education, publicly funded health services, and similar organisations share a structural characteristic that is profoundly favourable to incumbents: the individual making the software procurement decision is not symmetrically rewarded for savings and penalised for failures.

A public servant who selects a new AI-native platform that underperforms faces serious career consequences, including internal inquiries, negative press, public accountability and career termination.

The same public servant who retains a reliable incumbent and foregoes 30% cost savings faces no equivalent risk.

This asymmetry means that the rational decision for individual procurement officers in these environments is almost always to retain the incumbent, even when the economics of switching might appear compelling on paper.

Additionally, these sectors are motivated by service reliability rather than profitability, which further reduces sensitivity to price competition. This blunts the primary weapon of AI-native challengers.

## A Differentiated Risk Framework, Not a Binary Verdict

The ten factors above are not mutually exclusive, and their power compounds.

A SaaS business serving local government with a mission-critical, deeply integrated platform built on decades of proprietary operational data, in a heavily regulated sector, is almost entirely insulated from near-term AI disruption. This remains the case even if, in abstract terms, an AI-native challenger could theoretically replicate its features.

The challenger faces not one barrier but six or seven simultaneously, in an environment where the customer is institutionally predisposed to stay. In our view Technology One is the textbook example of such an incumbent.

Conversely, a SaaS business serving competitive, tech-forward customers with a generic horizontal product, shallow integrations, no proprietary data advantage, and easily portable customer data sits at the opposite end of the spectrum. This business is genuinely exposed, and the exposure is not a future risk. This risk is here now, as the recently released Claude Code product and the broader Claude Co-work AI ecosystem are demonstrating.

The task for investors is not to form a view on whether AI will disrupt SaaS — it will, selectively and unevenly. The bigger task is to map the specific characteristics of each business against a framework such as that proposed above, score

the moats honestly, and distinguish between those SaaS incumbents for whom AI is an existential threat and those for whom it is primarily a product enhancement opportunity.

By auditing a portfolio against these ten structural factors, an investor is better able to align their SaaS and tech investment strategy with higher “Structural Resilience”. The framework helps identify the tech and software companies best positioned to harness AI as a powerful margin-expander rather than being dismantled by it.

## AI Disruption Moat Audit (Our View)

### ASX-Listed Technology Companies — Structural Resistance to AI-Native SaaS Competition

#### AI Disruption Moat Audit — ASX Technology Companies

Scoring: 1 = Minimal moat | 2 = Low | 3 = Moderate | 4 = Strong | 5 = Very strong / structural

Company	Sector	Legacy Industry Customers	Deep Vertical Integration & Domain Expertise	Regulatory Burden	Monopolistic Industry Structure	Trust & Data Security	Tech Stack Integration Depth	Mission Critical Status	Network Effects & Data Pooling	Proprietary Non-Public Training Data	Public/NFP Procurement Asymmetry	Total (/50)
Technology One	Local Govt / ERP	5	5	5	4	5	5	5	5	5	5	49
Objective Corp	Govt Records Mgmt	5	5	5	4	5	4	5	3	4	5	45
Hansen Technologies	Utilities / Telco	4	5	4	5	4	5	5	4	4	2	42
Gentrack Group	Utilities / Airports	4	5	4	5	4	5	5	4	4	2	42
Iress	Financial Services	3	5	5	3	4	5	5	4	4	2	40
Nuix	Legal / Investigations	3	4	5	3	5	4	4	3	5	4	40
Wisetech Global	Logistics / Freight	3	5	4	3	4	5	5	4	5	1	39
Bravura Solutions	Superannuation / Life	3	5	5	3	4	5	5	3	4	2	39
Goria	Child Safety / EdTech	4	4	4	3	5	4	4	3	3	5	39
Hub24	Wealth Platform	3	4	5	3	4	4	4	4	4	2	37
Netwealth	Wealth Platform	3	4	5	3	4	4	4	4	4	2	37
PEXA Group	Property Settlement	4	5	2	2	5	3	5	3	2	4	35
Life360	Consumer Tech	1	4	3	3	5	2	3	5	5	3	34
Catapult Sports	Sports Analytics	3	5	2	2	3	3	3	4	4	3	32
Carsales	Auto Portals	2	4	2	3	2	3	3	5	4	1	29
REA Group	Property Portals	2	3	2	3	2	3	3	4	4	1	27
Siteminder	Hospitality Tech	2	3	1	2	2	3	3	4	4	1	25

Detailed assessments of two specific edge-case examples are provided in the Appendix below to further illustrate these risk factors.

**Disclosure:** The conclusions expressed herein were formulated by referencing independent market analysis from Ord Minnett and UBS. This article was produced with the assistance of AI in synthesizing these research inputs and drafting the final text.



**Victor Gomes**  
Principal and  
Portfolio Manager

## Appendix

### Technology One

**Local Government / ERP | Total Score: 49/50**

We believe Technology One is the standout most-protected business in this peer group. For this reason, we re-established a meaningful position after being allowed to buy it again in the last week of March when it was ejected from ASX Top 50 index. It scores at or near maximum on virtually every structural moat. Its core customer base consists of local governments, universities, and public sector bodies across Australia, New Zealand, and the UK — institutions defined by conservative procurement, asymmetric career risk for decision-makers, and near-zero price sensitivity to lower-cost software alternatives. The product is deeply mission-critical with finance, HR, asset management, student management (for universities), property and rating (for local councils) and procurement, all run on Technology One SaaS. It is vertically integrated into government workflows and supported by a compliance and audit capability that takes years to replicate. Proprietary data accumulated from decades of local government / university financials, planning records, and operational workflows cannot be trained on by an AI challenger. With a total moat score of 49/50, in our view Technology One represents the closest thing in Australian listed tech to a structurally AI-proof business.

Moat Factor	Score	Rationale
<b>Legacy Industry Customers</b>	5/5	Local governments, universities, and public utilities are among the slowest technology adopters globally. Procurement cycles run 12-36 months, IT teams are cautious, and switching decisions require council or board approval.
<b>Vertical Integration &amp; Domain Expertise</b>	5/5	Technology One's product suite covers the full ERP stack for government and higher education: finance, HR, asset management, planning, procurement, rating and student management. This represents 30+ years of customer and sector-specific workflow codification.
<b>Regulatory Burden</b>	5/5	Government customers operate under strict audit and compliance frameworks. Technology One must maintain accreditation across multiple regulatory regimes (APRA, state audit offices, TAFE compliance, etc.).
<b>Monopolistic Industry Structure</b>	4/5	Local government is not a competitive market. Councils face no pressure to reduce software costs to survive. The incentive to switch to save money is genuinely low.
<b>Trust &amp; Data Security</b>	5/5	Government and student data is highly sensitive. Vendor security accreditation (IRAP in Australia, UK-GDPR in the UK) is a prerequisite, not a differentiator.
<b>Tech Stack Integration Depth</b>	5/5	Technology One is typically the system of record — the ERP — meaning every upstream and downstream system connects through it. Integration switching cost is enormous.
<b>Mission Critical Status</b>	5/5	Financial management systems at councils and universities are as mission-critical as enterprise software gets. Downtime is politically and operationally unacceptable.
<b>Network Effects &amp; Data Pooling</b>	5/5	Technology One's benchmarking and shared services data across its large local government install base provides genuine peer-comparison value unavailable elsewhere.
<b>Proprietary Non-Public Training Data</b>	5/5	Decades of government and university financial transactions, asset records, planning approvals, and HR data is not available in the public domain.
<b>NFP Procurement Asymmetry</b>	5/5	Classic asymmetric incentive environment. Public servants bear significant career risk from failed software projects and minimal reward for cost optimisation. Strong status-quo bias.

## REA Group

### Property Portals | Total Score: 27/50

REA Group is a dominant marketplace business, not a traditional SaaS vendor, which changes the disruption calculus. Its network effect moat is genuinely strong — the combination of buyer intent data, listing depth, and agent advertising reach is self-reinforcing. Proprietary data (search behaviour, property transaction patterns, price movements at suburb level) is also meaningful. However, REA's customers (real estate agents) are not legacy industry operators, are moderately tech-forward, operate in a competitive industry, and face no regulatory burden that would slow adoption of new tools. The AI risk for REA is less about SaaS displacement and more about whether AI-native property portals (with better search, AI-generated listings, or agentic home-buying tools) could disintermediate REA's advertising model. Even in a less disruptive outcome where AI agentic tools direct traffic to REA's portal, the value of its depth product pricing differentiation becomes less justifiable. That risk is real and is only partially offset by the two-sided network effect.

Moat Factor	Score	Rationale
Legacy Industry Customers	2/5	Real estate agents are moderately tech-forward and competitive. They will adopt tools that generate leads faster, creating openness to new AI-native platforms.
Vertical Integration & Domain Expertise	3/5	REA's deep expertise in Australian property market data, advertising formats, and agent workflow tools is substantial and built over 25 years.
Regulatory Burden	2/5	Real estate is regulated but not heavily in terms of software compliance requirements. Barrier is low for new entrants on this dimension.
Monopolistic Industry Structure	3/5	REA effectively has a duopoly (with Domain) in Australian residential listings. Agent concentration in REA is high. However, the underlying industry is highly competitive, creating pressure to seek efficiency.
Trust & Data Security	2/5	Relatively low trust barrier. Property listing data is not sensitive in a security sense.
Tech Stack Integration Depth	3/5	REA integrates with CRM and property management tools but is not the system of record. Integration switching costs exist but are moderate.
Mission Critical Status	3/5	For an active agent, losing access to REA's listing platform would be a genuine business crisis. This creates strong retention but agents do have other options.
Network Effects & Data Pooling	4/5	The strongest moat: buyer intent data, listing depth, and agent advertising ROI data compound with scale. A new entrant cannot replicate this without a chicken-and-egg breakthrough.
Proprietary Non-Public Training Data	4/5	Suburb-level price trends, buyer search patterns, time-on-market data, and enquiry-to-transaction rates are genuinely proprietary and deeply valuable.
NFP Procurement Asymmetry	2/5	Private sector agents driven by commission. No procurement asymmetry — they will adopt whatever generates more business.

This material has been prepared by Eiger Capital Limited ABN 72 631 838 607 AFSL 516 751 (Eiger Capital). It is general information only and is not intended to provide you with financial advice or take into account your objectives, financial situation or needs. To the extent permitted by law, no liability is accepted for any loss or damage as a result of any reliance on this information.

Any projections are based on assumptions which we believe are reasonable but are subject to change and should not be relied upon.

Past performance is not a reliable indicator of future performance. Neither any particular rate of return nor capital invested are guaranteed.