Kitbag
ENTERPRISE EVOLVED

Why Not Australia

The Case for Sovereign AI Infrastructure in the Age of Geopolitical Disruption

Author	Dylan Nesbitt
Publisher	Kitbag Consulting
Initiative	Safe Harbour Initiative
Location	Canberra, Australia
Date	March 2026

TO DISCUSS FURTHER

Rod Robertson rod.robertson@kitbagconsulting.com.au · +61 411 702 785

Contents

Executive Summary

The Background

Every organisation that adopts enterprise software eventually discovers the same truth: the cost of leaving far exceeds the cost of staying. That is not a market failure. It is a business model. Salesforce, SAP, Oracle, ServiceNow, Palantir and a generation of System Integrators have built their scale by making institutional operations inseparable from their platforms. The logic governing how organisations function, who decides what and how authority is exercised migrates silently into vendor infrastructure. Once there, it is very difficult to retrieve.

The Argument

The digital sovereignty debate has focused on data residency: where information is stored. That question is real, but it is not the decisive one. What matters more is who controls the operating logic governing how that information is used: how decisions are made, how tasks are executed, how authority is exercised. That logic layer, the codified expression of how institutions function, is what must remain sovereign to the enterprise. Without it, no vendor can be exited on demand. Without a viable exit, there is no genuine choice. Palantir Technologies represents the most acute and most documented instance of this structural problem in Australia's national security environment. It is, however, the sharpest expression of a condition that ERP vendors, CRM platforms and their System Integrators have exploited for decades. AI compounds it significantly.

The Opportunity

Australia has a credible window to act on two fronts simultaneously: build the sovereign control layer that makes AI adoption safe and reversible and position itself as the stable, values-aligned jurisdiction that frontier AI companies and partner governments are actively seeking. Both goals follow from the same underlying condition. Retaining ownership of the operating logic that governs institutional decision-making is not a restriction on what technology Australia can use. It is the precondition for using any technology on Australia's own terms.

Introduction

Every major enterprise platform arrives with a genuine capability proposition and a contract structure that makes leaving progressively more expensive. Platforms marketed as Systems of Record[a] were built to store data and process transactions, and they do this well. What their vendors understood before most customers did was the commercial value of the layer below: each customisation an institution commissions, each integration it funds and each workflow it configures inside a platform quietly transfers operating logic from the institution into the vendor's environment. With every addition, the cost of leaving rises. After years of this dynamic, most institutions find their operating logic fragmented across vendor platforms, poorly documented and owned in practice by nobody except the systems that house it. What institutions have rarely been sold, and almost never built independently, is a true system of work[b]: a governed environment where that logic belongs to the institution itself and can be applied to whichever tools the institution chooses. For most organisations, that layer has never existed.

Systems of Record store data and process transactions. That is their function, and they perform it well. The more serious condition is that most institutions never built a system of work at all. Operating logic that should have lived in a sovereign, institution-controlled layer migrated instead into the customisation layers of their Systems of Record, encoded by vendors and held by them. The institution did not hand this over deliberately. It accumulated through each integration funded, each workflow configured, each module added. The vendor's continued cooperation is not a commercial convenience. It is a structural requirement.

A system of work requires owning operating logic holistically. Most institutions have never had this. Their operating logic is siloed inside vendor platforms, fragmented and extracted incrementally through each customisation they paid for.

The pattern predates AI and has been building for decades. SAP, Salesforce, Oracle, ServiceNow and a generation of System Integrators built their business models on exactly this dynamic. The exit cost is the product.

Artificial intelligence accelerates and deepens this condition. When AI systems execute tasks, analyse data and make recommendations, they do so according to logic embedded in the platforms they run on. An institution that has allowed its operating logic to migrate architecturally into a foreign vendor's platform has, in effect, allowed that vendor's logic to govern its AI outputs. The humans in the loop operate at the interface of a system whose underlying decisions they may not control and may not be able to inspect.

The decisive sovereignty question of the AI era is not where data is stored. It is who controls the operating logic governing how institutions function and exercise authority. Australia is currently answering that question by default, allowing operating logic to migrate into foreign-controlled platforms without a governed layer to preserve institutional independence. The pattern is visible across enterprise software broadly. In Australia's national security environment, one US company has achieved a level of

[a] System of Record: authoritative data store (transactions, records, assets). Holds outcomes; does not govern the process that produced them.
[b] System of Work: governed environment for planning, executing and tracking an institution's tasks, decisions and compliance obligations.

structural integration that makes the condition urgent and the stakes high.

The paper proceeds in five movements. It describes how vendor lock-in functions as deliberate commercial strategy. It examines how foreign law and geopolitical conditions create a second and distinct category of risk. It explains how AI agents compound both. It then interrogates Australia's specific situation, using Palantir Technologies as the primary case study, before drawing on international experience and setting out a structural response and three concrete recommendations.

The Lock-In Economy: Commercial Dependency by Design

The dependency that technology platforms and System Integrators create is not incidental to their business models; it drives them. The more deeply a vendor's platform integrates into institutional operations, the more costly substitution becomes and the more pricing power the vendor retains at every renewal.

This plays out across the enterprise software market in a consistent pattern. Salesforce becomes the System of Record for customer relationships, then the logic governing how those relationships are managed, escalated and resolved. SAP becomes the backbone of procurement and finance, then the rules governing how procurement decisions are made and approved. Oracle and ServiceNow embed into operations management, then define and customise the workflows that constitute institutional process. Each company offers genuine capability. Each also recognises that the longer they operate within an institution, the harder they become to remove.

> **Every customisation, integration and configuration an institution pays for transfers operating logic into the vendor's platform. That is not a side effect of enterprise software. It is the revenue model.**

The language used internally by these companies is instructive. The goal, as described in planning documents and strategy sessions, is to become the "operating system" of government or enterprise. The phrase is substantive, not aspirational. An operating system is what runs underneath everything else. You do not replace an operating system on demand; you plan for years, accept significant disruption and rebuild the processes that depended on it.

The costs of this dependency are real and well-documented. Australia's own experience with SAP offers a domestic proof. The Department of Defence's Enterprise Resource Planning program, launched in 2016 to consolidate hundreds of separate business applications into a single SAP S/4HANA system, was budgeted at between $1 and $2 billion with a target completion date of 2025. By 2024, the ANAO-audited program had blown out to $3.5 billion and extended to at least 2030: roughly three times the original cost, five years late.[1] The root cause identified in the First Principles Review that triggered the program was that Defence could not adequately document its own business processes.[2] The institution had lost track of its own operating logic. The vendor relationship became structurally irreversible not because the technology failed but because the institution no longer held an independent

record of what the technology was doing on its behalf.

The pattern extends beyond government. When Woolworths Australia transitioned from a thirty-year-old legacy system to SAP, store-level profit and loss reports that managers had relied on weekly could not be generated for nearly eighteen months.[3] The root cause was the same: the business had not documented its own operating logic before migration and institutional knowledge walked out the door with departing staff during a six-year transition. The vendor relationship was not the failure. The failure was the institution's inability to articulate its own operations independently of the vendor's system.

These cases share a common structure, and a common revelation. An institution adopts a capable platform. Integration deepens. The institutional understanding of its own operations atrophies because the platform handles it. When the institution eventually needs to exit or renegotiate, it does not merely discover that the vendor knows its operations better than it does. It discovers that it never owned its operating logic independently. The logic was always encoded in the vendor's layer: through customisations, integrations and configurations that belong to the platform, not the institution. The vendor did not take something. The institution built on someone else's land.

Vendors understand this and deliberately pursue this model, aided by their System Integrators. In September 2021, Bloomberg obtained internal emails from Palantir's UK regional head, Louis Mosley, with the subject line: "Buying our way in!"[4] The strategy he outlined involved acquiring smaller companies already serving the NHS to "take a lot of ground and take down a lot of political resistance." The company later called the language "regrettable." The underlying strategy continued. This is not unusual behaviour in the enterprise software market. It is the market. The question for governments and institutions is whether they have built the architecture to remain sovereign participants in it or whether they have become structurally dependent on its terms.

> *"Australia must reduce its reliance on imported technology and borrowed research."*
>
> Prime Minister Bob Hawke, 1990[5] – cited by the Ambitious Australia expert panel, December 2025, as 'more important than ever'[6]

Beyond Contract: How Foreign Law Reaches Australian Operations

Commercial lock-in and legal compulsion are two distinct categories of risk. Both operate through vendor dependency, but they differ in nature and origin. Commercial lock-in is a market dynamic: the product of deliberate business strategy that makes exit prohibitively expensive. Legal compulsion is structural: a feature of the international order that operates regardless of commercial intent or contractual arrangement. A vendor with good intentions and the most carefully drafted contract cannot override the laws of the jurisdiction in which it is domiciled.

The United States CLOUD Act compels American technology companies to produce data under their control regardless of where that data is stored.[7] Comparable extraterritorial legislation exists across multiple jurisdictions. The European Union, China and others each assert authority over companies operating within their legal frameworks that extends beyond national borders. Physical location does not resolve this exposure. “Sovereign cloud” infrastructure marketed specifically to government buyers addresses data residency, not legal jurisdiction. A US company running servers in Australian data centres remains subject to US legal orders. The relationship is structural, not contractual.

Legal compulsion is a distinct risk that shares the same common vulnerability: institutional operating logic held in platforms Australia does not control.

The corporate accountability problem compounds this. When Australian institutions need information or recourse from a multinational vendor, the accountability chain runs upward: to global headquarters, to foreign regulators, to legal obligations that may conflict directly with Australian interests. The PwC tax advice scandal illustrated the dynamic directly: when the Australian Senate and Taxation Office sought critical information, PwC International refused disclosure on the grounds that the material belonged to the global entity, not its Australian subsidiary. Multinational vendors answer ultimately to their global headquarters and foreign regulators. Under pressure, fiduciary duty flows outward. A domestically owned entity has no such conflict. Its accountability is anchored here.

The consequence for institutions that have allowed their operating logic to migrate into foreign-controlled platforms is qualitatively different from ordinary commercial risk. The exposure is three-fold. It is widespread: not bounded to a specific dataset or transaction, but extending across the entire operating logic the vendor has encoded: the rules, workflows, decision thresholds and governance structures that define how the institution functions. It is unquantifiable from the institution’s perspective: an organisation cannot audit what it cannot see, and the logic governing decisions made on its behalf may be inaccessible, subject to foreign direction, or modified without generating any visible change at the output level that a human reviews. It may also be unmitigatable by the time it surfaces: if a foreign legal order, policy shift or commercial decision has already shaped the operating logic, the problem may only become apparent when a decision reaches a human decision-maker, long after the underlying logic was set and the window for intervention has passed.

Three structural risk categories follow from this condition.[7] First, algorithmic bias: decision-support workflows that reflect the policy assumptions, legal frameworks and institutional priorities of the vendor’s home jurisdiction rather than Australian ones. Decisions shaped by foreign assumptions are not neutral. They embed the preferences and constraints of a different legal and political order into Australian institutional behaviour. Second, operational fragility: legal orders, sanctions or diplomatic disputes can disrupt platform access or constrain functionality at the platform level regardless of contractual protections and regardless of whether Australia is a party to or even aware of the underlying dispute. Third, loss of control: software updates, model changes and tooling constraints can be imposed

externally, without reference to Australian requirements or interests and without triggering any visible contractual breach. As AI systems become increasingly agentic, executing sequences of decisions autonomously, all three risks compound simultaneously and with less human visibility at each step.

The geopolitical environment has made this exposure more acute. Events of early 2026 demonstrated that technology relationships previously assumed to be stable can be fundamentally disrupted by political decisions in which Australia has no standing and no vote. Major technology companies, their engineers as much as their executives, have signalled publicly that unconditional cooperation with government demands is no longer guaranteed.

Australia is deepening its dependency on foreign AI platforms without building the governed operating layer that would make that dependency safe and reversible. Every sole-source contract, every contract variation, every security clearance that opens new doors to foreign platforms and every month without a sovereign ontology framework makes the eventual cost of correction higher and the strategic vulnerability deeper. The pattern Crikey's reporting makes visible: contract creep across Defence, AUSTRAC and intelligence agencies, each justified by claims that no alternative exists. That is exactly the dependency shadow that makes substitution prohibitively expensive even when the risks are well understood.

Allowing operating logic to migrate into foreign-controlled platforms creates risk that is simultaneously commercial and sovereign: the institution cannot audit what it cannot see, cannot exit what it cannot document, and cannot govern what it does not own.

> *"Flexibility without a defined destination is indistinguishable from drift, and drift is a luxury that nations do not have in a fast-closing technological gap."*
>
> Lowy Institute, March 2026[29]

AI at the Execution Layer: When Autonomous Systems Inherit Vendor Logic

The risks described in the preceding sections apply to any institutional dependency on foreign-controlled platforms. Artificial intelligence does not create these risks. It accelerates and deepens them in ways that fundamentally change the nature of the exposure, specifically by removing the human visibility that previously provided a partial buffer against both commercial and jurisdictional risk.

The progression is instructive. With traditional enterprise software, humans configure systems, inspect outputs and can override decisions with full visibility of the logic involved. With AI-assisted decision-making, humans remain nominally in the loop, but the reasoning producing recommendations becomes progressively opaque: the model's weights, training data and inference processes are not transparent to the user, and the institutional logic embedded in the platform shapes outputs in ways that may not be visible at the interface. With agentic AI (systems that plan, act and execute sequences of decisions autonomously), humans review outcomes, not the reasoning, data ingestion or operating logic that produced them. The institutional visibility that previously buffered both commercial and jurisdictional risk disappears.

> **Agentic AI removes the human visibility that previously mitigated both commercial and jurisdictional risk. Every autonomous decision executes according to logic the institution may not own.**

The practical consequence is direct. An institution whose operating logic has migrated into a foreign-controlled platform is already exposed to the risks described in the preceding sections. When AI agents operate against that logic, executing institutional tasks autonomously, the institution has effectively outsourced not just its data and workflows but the decision-making itself. The logic governing those decisions was not written by the institution, is not inspectable in real time and cannot be overridden without abandoning the system entirely. The humans who review the output may have no visibility of what the system did to produce it, what data it drew on, or whether that data or the logic applied to it was subject to foreign direction, modification or constraint.

This is not speculative. The Pentagon-Anthropic dispute, examined in detail in the following section, revealed that the most powerful military on earth had embedded a single AI vendor so deeply into classified operations that disentangling it required a six-month phase-out. The operating logic governing intelligence analysis, operational planning and reportedly near-autonomous operational support had migrated to the vendor's platform without a governed layer to preserve institutional independence. The condition became visible not through institutional review or audit, but through a commercial dispute. That is precisely the detection failure this section describes.

For Australia, the implication is concrete. As government agencies adopt AI tools contract by contract with no sovereign ontology framework in place, each adoption embeds more operating logic into platforms Australia does not control. Each autonomous action those systems take deepens the dependency. Each decision that surfaces to a human decision-maker may have been shaped by logic that is foreign-governed, unauditable and, in the event of a commercial or geopolitical disruption, potentially inaccessible. The window to build the governed layer that makes this exposure manageable narrows with each contract signed.

The Palantir Case: The Canary in the Coalmine

Palantir Technologies represents the most acute instance of the structural problems described above in Australia's national security environment. The company offers genuinely capable analytical tools. It has also pursued a strategy of institutional integration that goes further, faster and into more sensitive environments than any previous enterprise vendor in this domain.

Australia's Department of Defence awarded Palantir a $7.6 million contract for its Cyber Warfare Division in February 2026, without competitive tender.[8], [9] Defence justified bypassing the open market by claiming no other software could match Palantir's capabilities. The contract brings Australia's total spend with the company to more than $26 million since 2013.

Palantir solved the fragmentation problem that Systems of Record never could: it unified operating logic across silos. Then it put that logic in a proprietary black box. Consolidated dependency is harder to exit than fragmented dependency.

The dollar figures understate the nature of the arrangement. In March 2026, Crikey obtained a copy of a 2024 Palantir-Defence contract for Industrial Intelligence Capability, the first time the terms of Palantir's work for the Australian government had been made public.[10] The contract reveals Palantir specialists embedded permanently on-site at Defence; cybersecurity notification obligations triggered only by confirmed breaches rather than suspected ones; penetration testing of Australian Defence systems requiring Palantir's consent; and provisions allowing Palantir to train models on Defence analyst behaviour. A senior Defence official, speaking anonymously, put the structural concern directly: "Why aren't Australians checking Australians?"

The pattern extends beyond Defence. AUSTRAC awarded Palantir an $8.1 million contract for data analytics in 2023, then executed five contract variations in the twelve months that followed, pushing the total past $12 million. Further variations followed in early 2026. The Australian Criminal Intelligence Commission has followed a similar trajectory. In November 2025, Palantir received formal Australian government security assessment at the Protected level, explicitly opening the door for a broader range of Commonwealth agencies to access its platforms.[11] The company described Australia as "an important market."

A financial arrangement sits alongside these contracts. Australia's Future Fund holds 498,339 shares in Palantir Technologies, valued at $103.6 million as at June 2025.[12] When Senator Pocock questioned the Fund's chief corporate affairs officer at Senate Estimates in February 2026, asking whether it would divest from companies "profiting from surveillance, from weapons and from human suffering," the officer could not commit to divestment.[13] The sovereign wealth fund that manages money on behalf of the Australian government is, in effect, a shareholder in the company whose operating logic now governs how parts of that same government function.

Parliamentary concern is real and cross-party. Senator David Shoebridge has described Palantir as "a global surveillance empire" and the Future Fund's continued investment as "a deep and fundamental breach of public trust."[14] Digital Rights Watch submitted directly to Finance Minister Gallagher that the

government must “prioritise our human rights, privacy and digital sovereignty over Palantir’s dystopian surveillance profiteering.” These are not fringe positions. They reflect a documented gap between what government procurement is producing and what Australian institutions should be able to account for.

> *“Palantir is a company based in the USA, where there is a possibility that sensitive data could be accessed by the American government and intelligence services.”*
>
> Swiss Army internal risk assessment, 2024[15]

Palantir’s own product literature illuminates both the genuine capability and the structural trap. Its Foundry system is built around what the company calls an “ontology”: a structured model of an organisation’s entities, relationships and decision logic. Foundry achieves something no System of Record ever managed: it extracts operating logic from fragmented silos and unifies it into a coherent, queryable layer. That is its real capability, and it explains why institutions find it genuinely useful. The trap is that the unified operating logic lives inside Palantir’s proprietary platform. The institution has not gained ownership of its operating logic; it has traded fragmented vendor dependency for consolidated vendor dependency, and the consolidated version is considerably harder to exit precisely because Palantir succeeded where others failed. An institution that has given Palantir its operating logic has, in the most complete sense available, surrendered its system of work to a foreign-governed black box. A sovereign operating ontology achieves what Palantir offers under the opposite governance arrangement: the operating logic belongs to the institution, is auditable by the institution, and can be applied to whichever AI or analytics tools the institution chooses, including replacing them entirely without losing the institutional knowledge encoded in the layer beneath.

The Pentagon–Anthropic Precedent

The structural argument in this paper is not theoretical. In February 2026, the United States government provided a large-scale proof of what happens when institutional operations become inseparable from a single vendor’s product.

On 28 February 2026, the US President ordered every federal agency to immediately cease use of Anthropic, the company behind Claude, at the time the world’s most widely deployed enterprise AI model and the only AI system on the Pentagon’s classified networks. The Pentagon simultaneously designated Anthropic a national security supply chain risk, a classification previously reserved for foreign adversaries. Anthropic’s offence was refusing to permit its models to be used for mass domestic surveillance of American citizens and fully autonomous lethal weapons without human oversight.

> **Six months to exit a single AI vendor. The condition that made this necessary was built contract by contract, with no sovereign layer in place. The cost was operational paralysis at the highest level of government.**

The six-month phase-out period announced alongside the ban said more than the ban itself. The most powerful military on earth, which had used Claude in intelligence analysis, operational planning and

reportedly in the operation to capture Venezuelan leader Nicolas Maduro, could not exit a single vendor in less than half a year. Not because the model was technically irreplaceable (OpenAI signed a replacement deal within hours),[16] but because the operational integration had been built without a governed layer between mission and tool. When interests diverged, the institution discovered it had no architecture for the transition. That is a governance failure, not a procurement failure.

The OpenAI deal that followed illustrates the same vulnerability from a different angle. Where Anthropic had sought explicit contractual prohibitions on specific misuses, OpenAI's agreement relied on assumptions that federal agencies would not break existing law. Analysts at MIT Technology Review described it as "softer legal language rather than freestanding contractual rights."[17] OpenAI's own chief executive subsequently acknowledged the deal had been rushed. A different company did not solve the dependency problem. Institutional dependency on any single vendor, whatever its identity, is itself the structural vulnerability.

Dean Ball, a senior fellow at the Foundation for American Innovation who served briefly as a Trump administration AI policy advisor, described the Pentagon's actions as "attempted corporate murder" and "a psychotic power grab" that sent a disqualifying message to any business considering operating in the United States.[18] His objection was structural, not partisan: the government had demonstrated it would destroy a company that refused to remove the safeguards its platform depended on, rather than building the institutional architecture that would have rendered the dispute operationally irrelevant.

As of this paper's publication, Anthropic is suing the Department of Defense, challenging the supply chain risk designation as "unprecedented and unlawful."[19] A federal court hearing is scheduled for 24 March 2026. Whatever the legal outcome, the operational lesson is already established: when a government has no sovereign control plane between its mission and a vendor's product, it has no exit.

The International Reckoning

Australia is not alone in confronting this question. Countries that have examined the structural terms of vendor dependency in sensitive environments have, in most cases, moved to limit or exit it. Countries that did not move early enough are now managing the consequences of dependencies they cannot reverse on demand.

Switzerland. In 2024, the Swiss Army conducted a formal twenty-page risk assessment of Palantir and concluded the company posed unacceptable risks to national data control.[15] The assessment found that Palantir's system would require company specialists permanently on-site, limiting Switzerland's ability to act independently in a crisis. Swiss authorities rejected Palantir on those grounds despite being pitched by the company's senior UK representatives. They did not question the technology's capability. They questioned whether a sovereign state could accept structural dependency on a foreign vendor for the logic layer governing its defence operations. Palantir was rejected at least nine times by Swiss federal agencies across seven years of attempts.

Switzerland rejected Palantir nine times across seven years. Denmark is now seeking an exit. The UK is reckoning with the cost of not having done so earlier.

Germany. Sinan Selen, the head of Germany's domestic intelligence service, told a Berlin conference in December 2025 that Europe "must be able to generate alternatives" to CIA-backed platforms like Palantir, "taking into account geostrategic considerations."[20] He described software selection as requiring a three-factor test: security contribution, performance and whether it is "geostrategically correct." A German civil rights researcher described Palantir as "the AI arms dealer of the 21st century,"[21] warning of the geopolitical and legal risks in handing sensitive police data to a US company subject to conflicting foreign laws.

Denmark. Following Trump's escalating threats over Greenland, Danish intelligence services are actively seeking a replacement for Palantir's data processing platform.[22] The Danish Defence Intelligence Service published an assessment in late 2025 that, for the first time in its history, named the United States as a potential security concern, describing how Washington uses "its economic and technological strength as a tool of power, also toward allies and partners." The implication for technology dependency on US-owned platforms is direct.

The United Kingdom. The UK is the cautionary case. The MoD awarded Palantir a 240 million pound no-tender contract in December 2025. Palantir now holds over 670 million pounds in UK government contracts spanning defence operations, NHS patient data, police intelligence databases and nuclear weapons management. Liberal Democrat MP Martin Wrigley raised this in parliament:[23] "Will the government look into ensuring that Palantir is not a single point of failure in our critical systems, in the health service, defence, the Cabinet Office and now the police?" Two serving Ministry of Defence systems engineers broke cover in March 2026 to warn that Palantir poses "a national security threat" to Britain, saying ministerial assurances about data ownership "are ignorant and miss the point entirely."[24] Their concern is what intelligence professionals call the mosaic effect: Palantir's combined access across defence, health, policing and critical infrastructure creates a complete operational picture of the sovereign state. Individual contractual protections over specific datasets do not address this. The picture is the problem.

The pattern is consistent. Countries that examined the structural terms carefully declined or retreated. Countries that moved fast are now managing dependencies they cannot reverse on demand. Australia is at the decision point Switzerland reached first and got right.

The Two-Layer Solution

Building domestic AI models, excluding foreign technology or retreating from the global ecosystem would trade one problem for another. Institutions that cannot access the best available tools will be outperformed by those that can. Sovereignty purchased at the price of capability is not a viable settlement.

The European Union has moved toward exactly this kind of trade-off. Its digital sovereignty agenda, pursued through the Cloud and AI Development Act, GAIA-X and related initiatives, aims to reduce dependence on US technology by building European alternatives. The Atlantic Council's analysis of this effort is instructive: the GAIA-X initiative produced largely a series of standards and labels rather than a transformation of the commercial landscape. Mario Draghi's 2024 report on the EU single market "effectively conceded defeat in this area of endeavour."[25] Sovereignty through exclusion has not worked. Australia's approach should be different.

Owning a sovereign operating ontology means owning institutional logic that, for most organisations, has never been truly theirs. That ownership is what makes any vendor: AI model, analytics platform, or System of Record, genuinely replaceable.

The resolution lies in an architectural distinction between two layers that must be governed differently.

The first is execution infrastructure: the models, platforms, compute and tooling that perform work. This layer can and should be globally shared. Australian institutions should have access to the best AI capabilities available, wherever they originate. Attempting to replicate frontier model development domestically is neither feasible nor necessary.

The second is the operating ontology: the formal, machine-readable representation of how an institution functions, covering its entities, relationships, rules, decision points, responsibilities and governance structures. This is the codified expression of institutional will. It translates parliamentary intent, through regulation and policy, into executable action. This layer must be sovereign. For most institutions, building it would mean owning their operating logic for the first time; not recovering something lost, but constructing something that was never fully theirs. A sovereign system of work is what that ownership looks like in practice: an environment through which the institution's operating logic is maintained, governed and applied by the institution itself, making it an institutional asset rather than a vendor's leverage point.[26]

The distinction matters for a reason that goes beyond governance. Two organisations doing identical work will have materially different operating logic: different approval thresholds, different escalation

paths, different compliance interpretations, different efficiency heuristics built up over years of institutional experience. That accumulated intelligence is what separates institutions that perform from those that merely process. In the private sector, it constitutes competitive advantage. In the public sector, it constitutes the uniquely Australian expression of democratic authority: the specific way this institution, under this legal framework, exercises the judgement Parliament intended. Ceding it to a foreign-governed platform is a decision about who holds the crown jewels, not a technology procurement call.

As AI capability grows, the translation layers between human intent and machine execution are thinning. Models are becoming more capable, more interchangeable, more commoditised. What remains permanent is the codified understanding of what an institution does, why it does it and how it exercises judgement. An institution that owns its operating ontology can replace any AI tool as a straightforward procurement decision. An institution that has ceded its operating logic to a vendor cannot replace that vendor without replacing its own institutional memory.

The Pentagon found itself in that condition. Not because Claude was irreplaceable as a model (OpenAI signed a deal within hours), but because the operational integration had been built without a governed layer between mission and tool. Had that sovereign layer existed, the transition would have been a procurement event rather than an operational crisis.

Genuine sovereignty in the AI era is, at its core, about market leverage. Countries should be able to use whatever frontier AI tools best serve their goals, regardless of where those tools originate. That freedom requires the structural capacity to switch. Buyers who cannot replace a vendor have no bargaining power, no accountability mechanism and no sovereignty in any operational sense. The operating ontology is what creates and preserves that capacity.

The Australian government's own AI planning acknowledges pieces of this. The APS AI Plan introduces GovAI Chat, a secure, government-controlled generative AI tool operating within Australian government infrastructure with data remaining in-country.[27] A meaningful step. It addresses data sovereignty, not operational sovereignty. It says nothing about who controls the decision logic, the workflow rules or the governance structures governing how Australian institutions function. The answer to that question is the architecture this paper proposes.

In December 2025, the Ambitious Australia expert panel delivered its Strategic Examination of R&D; to three federal ministers.[6] The panel, chaired by Robyn Denholm of Tesla and including former Chief Scientist Ian Chubb, found that "widespread views that Australia can simply purchase innovations from overseas will yield a bleak Australian economy and greater sovereign risk," and recommended an "if not, why not" principle for Australian procurement. The report noted that 23 per cent of Commonwealth procurement already goes offshore. It is serious work on research, capital and workforce policy. It does not address the operating logic layer: the question of who governs how Australian institutions actually function when AI executes their decisions. These two arguments reinforce each other. Together they represent more institutional weight behind the sovereign technology case than has existed at any point in Australia's policy history.

> *"Unless Australia shapes the rules of the game, we will end up as a passive technology taker, chained to systems we neither understand nor control. AI is not born neutral. If Australia adopts AI built to foreign specifications, we also import those embedded choices about privacy, bias, autonomy, and control."*

Lowy Institute, 2025[28]

The Democratic Imperative

Anthropic drew two lines, no mass surveillance of citizens and no fully autonomous lethal weapons, and held them at extraordinary commercial cost. The company was valued at $380 billion before the executive action, held the Pentagon's only classified-network AI contract and chose principle over compliance. Every democracy will face a version of that negotiation. The question is whether governments face it having built sovereign institutional infrastructure, with the capacity to audit tools, explain decision logic and replace vendors, or mid-dependency, unable to remove a tool their operations cannot function without.

Institutions that cannot audit their AI tools, explain their decision logic or replace a vendor without operational paralysis are not functioning as democracies. They are functioning as dependencies. The sovereign control plane is the expression of democratic values applied to AI governance. Without it, oversight is a performance. With it, accountability is structural.

Sovereign AI infrastructure is the structural expression of democratic accountability. Where it is absent, oversight is performative. Where it exists, accountability is built into architecture.

This challenge extends well beyond Australia's immediate region. Governments on every continent are grappling with the same structural choice: adopt powerful AI tools built by foreign companies under foreign legal regimes, or forgo the capability. Nations across Southeast Asia, the Pacific, Europe and the Global South are caught between US and Chinese AI vendors, both of whom come with strategic strings attached. The Pentagon-Anthropic dispute demonstrated what US dependency looks like when interests diverge. Chinese AI platforms carry their own well-documented governance and surveillance implications. There is space for a third approach: a governance model that demonstrates how sovereign AI infrastructure can be built within a democratic framework without forfeiting capability.

Australia, as a stable democracy, AUKUS member and respected middle power with institutional relationships across the region and beyond, is a credible candidate to develop and demonstrate that model. Not as a technology superpower competing on capability, but as the jurisdiction that shows how the governance problem is solved.

Switzerland did not build the world's largest banks. It built the world's most trusted legal and institutional framework and capital came to it. Every nation that wanted to protect assets across jurisdictions eventually needed what Switzerland had designed. Today, every nation adopting AI faces the same sovereign governance question. The country that solves it first, demonstrating in practice how to

harness frontier AI tools while retaining democratic control over the logic layer governing their use, will have built something genuinely exportable. Not a product. A proof. The legal conditions, the policy architecture, the institutional frameworks that make sovereign AI governance operational rather than aspirational.

Australia has the standing to be that country. It has the democratic institutions, the strategic relationships, the policy community and the geopolitical incentive. What it needs is the decision to act on that position before the window that makes it credible closes.

Three Concrete Recommendations

Contract clauses, ANAO audit findings, the testimony of serving engineers across three continents and the geopolitical disruptions of the past three months give the risks in this paper measurable form. The window for Australia to act is measured in months, not years. Three actions would materially change Australia’s position.

Three actions would materially change Australia's sovereign risk profile. None require abandoning foreign technology. All require building the layer that makes foreign technology replaceable.

RECOMMENDATION ONE

Establish a Sovereign Infrastructure Framework

The Department of Finance should mandate that all Commonwealth agencies and critical infrastructure operators maintain sovereign operating ontologies: formal, machine-readable models of institutional operations that are Australian-owned, Australian-governed and subject only to Australian law. This does not require abandoning foreign Systems of Record. It requires building the governed layer, a System of Work, between institutional mission and tool that the Pentagon did not have when it needed it most. The framework should draw a clear line between execution infrastructure, which can be globally sourced, and the operating ontology, which must be domestically controlled. The APS AI Plan’s GovAI initiative is a foundation. This framework is the architecture that should sit above it.

RECOMMENDATION TWO

Commission an Independent Review of Vendor Dependencies

Australia should conduct a strategic review of current vendor dependencies across Defence, intelligence, law enforcement and critical infrastructure, modelled on the Swiss Army's 2024 Palantir risk assessment. The review should assess the sovereign risk profile of existing contracts, identify where operating logic has migrated into foreign-controlled platforms, assess whether the IRAP clearance process adequately addresses jurisdictional risk rather than technical security alone and recommend managed transition strategies with defined timeframes. The question is not whether to exit relationships with capable foreign technology companies. The question is whether Australia has the architecture to do so if it ever needs to.

RECOMMENDATION THREE

Position Australia as a Trusted Governance Partner

Australia should engage proactively with companies seeking stable, values-aligned jurisdictions and develop sovereign AI governance frameworks that can be offered to partners globally navigating the space between US and Chinese platforms. This means arriving at AUKUS technology-sharing mechanisms, multilateral forums and bilateral technology discussions with a clear position on what sovereign AI infrastructure looks like and why it makes partnerships more durable, not less.

Conclusion

The Pentagon-Anthropic dispute was not an anomaly. It was a preview of a structural condition that every institution adopting AI will eventually face: the moment when the interests of the tool provider and the interests of the institution diverge and the institution discovers whether it built the architecture to manage that divergence or merely hoped it would never arrive.

The structural problem is not new. Enterprise software vendors and their System Integrators have made institutional operating logic inseparable from their platforms for decades. Artificial intelligence accelerates and deepens this condition while removing the human visibility that previously buffered the risk. In Australia's national security environment, one company has achieved a level of structural integration that makes the condition acute. The response is to build the sovereign layer that makes those tools genuinely replaceable.

Australia can learn from the preview the Pentagon provided rather than replicate its outcome. The operating ontology that sits between institutional purpose and AI execution is not optional infrastructure. It is the minimum architecture for safe, scalable, sovereign AI: what makes tools genuinely replaceable, decisions auditable and democratic governance compatible with AI capability rather than subordinate to it.

Switzerland understood this before a crisis forced the question. Denmark is understanding it now. The UK is understanding it afterwards: 670 million pounds in contracts that cannot be easily unwound, live parliamentary investigations across health, defence and policing and serving systems engineers warning publicly of a national security threat.

> *"A nation's ability to adopt, replace and govern AI tools without institutional paralysis is now the defining measure of sovereign resilience."*

Australia can build that resilience deliberately, or discover its absence the hard way. The window is open. It will not remain so indefinitely.

TO DISCUSS FURTHER

Rod Robertson

rod.robertson@kitbagconsulting.com.au · +61 411 702 785

Sources and Notes

1. The Canberra Times, 'Defence ERP overhaul blows out to $3.5b in latest IT bungle,' 17 September 2024; InnovationAus, 'Defence ERP blows out as switchover looms,' 19 September 2024

2. ANAO Performance Audit, 'Defence's Administration of Enabling Services: Enterprise Resource Planning Program: Tranche 1,' 2021-22; First Principles Review, Department of Defence, 2015

3. CIO Magazine / ARN, 'Famous ERP disasters,' citing Woolworths Australia SAP transition; commentary from implementation consultants on process documentation failure

4. Bloomberg, 'Palantir had plan to crack UK health system: Buying Our Way In,' 30 September 2022

5. Prime Minister Bob Hawke, speech on Australian R&D;, 1990 (as cited in Ambitious Australia, December 2025)

6. Ambitious Australia: Strategic Examination of R&D; Final Report, independent expert panel to Ministers Ayres, Chalmers and Clare, December 2025

7. Kitbag Consulting, 'The Sovereign Synapse: Why Australia's Operating Logic Must Remain in Australian Hands,' February 2026, kitbagconsulting.com.au/news. Sources for CLOUD Act: 18 U.S.C. §2523 (Clarifying Lawful Overseas Use of Data Act, 2018). Sources for PwC corporate veil: Senate Finance Committee proceedings, 2023; AFR and The Australian, reporting on PwC International's refusal to disclose documents to the Australian Senate

8. Crikey (Cam Wilson), 'Defence signs biggest ever contract with Palantir for Cyber Warfare Division,' 17 February 2026

9. InnovationAus, 'Palantir handed $7.6m contract to map Australia's Defence industry,' 18 February 2026

10. Crikey (Cam Wilson), 'Palantir embedding staff in Defence and mining Australian data,' 9 March 2026

11. Michael West Media, 'Palantir secures top Australian security clearance,' January 2026

12. Honi Soit, 'Australia's $100 million investment in Palantir,' February 2026; Future Fund Annual Report 2025

13. Senate Estimates, February 2026; Startup Daily, 'Australia's Future Fund invested $103 million in Palantir,' February 2026

14. Senator David Shoebridge, public statements and Senate Estimates, February 2026

15. Swiss Army internal risk assessment of Palantir, 2024, as reported in Swiss magazine Republik and The Guardian

16. NPR, 'OpenAI announces Pentagon deal after Trump bans Anthropic,' 28 February 2026

17. MIT Technology Review, ‘OpenAI’s compromise with the Pentagon is what Anthropic feared,’ 2 March 2026

18. Fortune, ‘The fight between Anthropic and the Pentagon raises crucial questions,’ March 2026 (Dean Ball characterisation); Foundation for American Innovation, public commentary

19. CNBC and Reuters, multiple reports on Anthropic lawsuit and Pentagon designation, March 2026

20. Reuters and Heise Online, Sinan Selen remarks at BfV symposium, Berlin, December 2025

21. The Local (Germany), ‘What is Palantir’s Gotham software and why do German police want it,’ August 2025

22. Progressive International, quoting Danish intelligence services seeking Palantir replacement, February 2026

23. openDemocracy, ‘The great Ministry of Defence-to-Palantir pipeline,’ February 2026; UK Hansard, ‘Ministry of Defence: Palantir Contracts,’ 10 February 2026 (Wrigley, Lewis, Cooper)

24. The Nerve / AOAV, ‘Britain’s Palantir problem,’ 16 March 2026 (MoD insiders, mosaic effect quotation)

25. Atlantic Council, ‘Digital Sovereignty: Europe’s Declaration of Independence?’ February 2026; referencing Draghi 2024 EU competitiveness report

26. Kitbag Consulting, ‘Why Operations Ontologies Are the Critical Foundation for AI-Driven Enterprise Evolution,’ March 2026, kitbagconsulting.com.au/news

27. Department of Finance, APS AI Plan 2025, digital.gov.au

28. Lowy Institute, ‘Australia’s AI choice: Standards setter or technology taker,’ 2025

29. Lowy Institute, ‘Australia bets on old laws to manage new AI risks,’ March 2026

This paper is a Safe Harbour Initiative on sovereign AI infrastructure in Australia · safeharbourinitiative.com.au · Kitbag Consulting · kitbagconsulting.com.au · March 2026