SUBMISSION

GETTING IT RIGHT: BUILDING AI INFRASTRUCTURE THAT WORKS FOR AUSTRALIA

Capturing operating logic as the precondition for Australia's AI dividend



DOLIUM

Kitbag Consulting
KB Technology Pty Ltd | DOLIUM

Submission to the Department of the Prime Minister and Cabinet
September 2026

SUBMISSION

Capturing operating logic as the precondition for Australia's AI dividend

Submitted by: Rod Robertson, CEO, Kitbag Consulting; Director, KB Technology Pty Ltd (DOLIUM)

In response to: Department of the Prime Minister and Cabinet consultation paper, September 2026. Submissions close 5pm AEDT, 9 October 2026.

Contents

Executive Summary 3
Getting it right: Building AI infrastructure that works for Australia 5
1. Frontier is not required for BAU and AI dividend should not be contingent on it 5
2. Token efficiency, captured operating logic and the energy and cost objectives 7
3. Transparency, repeatability and operable safety disclosure 9
4. The commoditisation and replacement risk and third party access 11
5. Aggregation, sovereignty and the definition the standards should adopt 12
6. What operating logic capture requires in practice 12
Conclusion 13
Recommendations 15

Executive Summary

Kitbag Consulting and KB Technology work with Australian public and private sector entities on digital initiatives and AI governance. KB Technology has developed DOLIUM, an Australian owned, IRAP assessed System of Work listed on the Digital Transformation Agency Software Marketplace Panel. This submission draws on that engagement, on submissions already made to the Joint Select Committee on Artificial Intelligence, the Digital Transformation Agency and the Department of Home Affairs and on direct briefings to the Department of Finance and other agencies on the same subject.[1]

The consultation paper addresses two related but distinct aspects of Australia's AI capability. The first is physical: the infrastructure, energy, water, land, community and workforce requirements associated with large data centres. The second concerns conditions for frontier AI training in Australia, including national security and safety, local capability, skills, research and innovation, and support for Australian creators and rightsholders. Both are necessary. Neither however, directly addresses the organisational operating layer through which AI produces economic and governance outcomes.

This submission does not seek to duplicate those requirements or to address every question raised by the consultation. It identifies a cross-cutting omission: the operating logic of the organisations that will use this infrastructure and these models, meaning the decision rights, policies, processes and institutional knowledge that govern how government and enterprise function.

This submission argues that capturing this layer, in a sovereign Australian governed structure is a precondition for realising the paper's own stated objectives, not an adjacent policy interest.

Six findings follow, each tied to a specific part of the paper.

- First, the broader Australian AI productivity dividend should not be made contingent on attracting frontier training activity. The consultation correctly identifies genuine strategic, economic and national security benefits from frontier training in Australia. However, the substantial majority of routine government and enterprise AI use does not require frontier reasoning. This creates an opportunity to pursue AI productivity through smaller, open-weight or otherwise substitutable models operating against structured organisational context, while separately pursuing the national benefits of frontier training.
- Second, capturing operating logic can reduce the token, compute and energy footprint of AI use under either a hosted or a self-hosted deployment model. This creates a demand-side efficiency opportunity that complements the consultation's proposed supply-side requirements for energy, water and cost management at large data centres.

- Third, a structured operating logic layer can make the paper's proposed safety and security objectives more traceable in practice. Model-level incident reporting can explain what went wrong within a model. It does not necessarily establish what the model was authorised to do inside an organisation, what process it was executing, what decision authority applied or which human was accountable for the resulting action.
- Fourth, without that layer government and enterprise are exposed to two related risks that are not adequately addressed by model-level disclosure alone: the progressive embedding of organisational operating logic into vendor platforms and the potential for foreign third parties to gain structural visibility into how an organisation operates.
- Fifth, sovereignty should be considered in the eventual standards as including effective Australian control over operating logic, rather than being defined solely by data residency or the physical jurisdiction in which a model is hosted or trained. Data residency is necessary in some circumstances. It is not by itself sufficient to establish operational sovereignty.
- Sixth, the cost of getting this sequence wrong can be material. The Commonwealth and Australian jurisdictions have already experienced the difficulty and cost associated with major enterprise platforms that become deeply coupled to organisational requirements. Those examples do not establish that architecture was the sole cause of those outcomes, but they demonstrate why operating requirements should be represented independently before major platform investments are made.

The recommendations at the close of this submission set out how the AI standards can recognise each of these findings.

Getting it right: Building AI infrastructure that works for Australia

1. Frontier is not required for BAU and AI dividend should not be contingent on it

The paper is right that frontier training investment carries genuine national security, research and economic benefit and that Australia should compete for it on its own terms. However, it would be a mistake to let the case for that investment imply that Australia's broader AI dividend depends on it.

In Kitbag Consulting's assessment, drawn from engagement across government and enterprise AI deployments, the substantial majority of routine AI use does not require frontier reasoning. It involves retrieving a policy, summarising a submission, routing a request to the correct delegate or checking a proposal against an approval pathway. That work is well served by smaller, open-weight models operating inside a structured operating context, because a modest model given the organisation's actual decision logic can be more useful for a defined business process than a larger model given none.

This matters for two reasons directly connected to the paper's own concerns.

- The first is economic. If the substantial majority of AI derived productivity is available through smaller models running on modest, domestically manageable compute, then Australia's AI driven productivity growth, the benefit the paper cites data centre investment as delivering, is not conditional on whether frontier training activity is won here or on any single frontier provider's continued willingness to operate under Australia's proposed conditions.
- The second is continuity. Frontier model access has already proven to be a function of geopolitics rather than a stable utility. On 12 June 2026, the US government applied export controls to Anthropic's Claude Fable 5 and Claude Mythos 5 models, requiring Anthropic to restrict access by foreign nationals. Because the company could not reliably verify nationality in real time, Anthropic suspended access to both models for all users. The controls were lifted on 30 June and access was restored from 1 July.[2]

The significance of this event for Australia is not the particular models involved. It is that access to a frontier model can be altered by the legal and geopolitical circumstances of its provider, independently of where an Australian customer or infrastructure happens to be located.

No Australian contract, subsidiary structure or physical hosting arrangement can be assumed to eliminate that exposure. An Australian deployment may therefore remain dependent on the continued ability and willingness of an offshore provider to make its model available.

Onshore training delivers real and separate benefits, including local jobs, energy investment, research capability, skills and the opportunity to influence the development of AI safety and security standards. This submission supports those objectives. It does not, however, follow that onshore training alone resolves continuity or sovereignty risk.

The paper itself recognises that Australia should develop AI infrastructure and capability "on terms that work for Australians" and that the benefits of investment need to remain in Australia's national interest.

The relevant architectural question is therefore not whether Australia should host frontier AI. It is whether Australian organisations can continue to operate when the model behind an AI-enabled process changes, becomes unavailable or is subject to a foreign provider's legal obligations.

Where the operating logic resides within the enterprise rather than inside a vendor platform, an AI agent is a role the entity assigns, not a capability welded to one provider. When a given model becomes unavailable, whether by export control, licensing change or commercial withdrawal, the entity can, subject to the technical capabilities of the replacement model, re-point the governed role to another model without rebuilding the decision logic, delegations and process context that made the agent useful in the first place.

The role can therefore continue to run against the same structured operating context regardless of which model sits behind it. The continuity benefit comes from separating organisational operating logic from model capability.

Business continuity under this model is not a property of any single vendor's resilience. It is a property of how deliberately the entity has separated its operating logic from whichever model happens to be running it on a given day.

2. Token efficiency, captured operating logic and the energy and cost objectives

Part 2 of the consultation paper proposes mandatory requirements for large data centres to contribute to Australia's energy transition, minimise water use, avoid shifting costs to consumers and communities and contribute to skills and training.

These are proportionate and necessary requirements for the physical layer. They do not, however, fully address demand-side waste, meaning inference activity that consumes energy and compute without producing proportionate useful output. The consultation therefore has an opportunity to consider not only the efficiency of the infrastructure supplying compute, but the efficiency of the workloads consuming it.

Where an organisation's operating logic has been captured and made machine readable, KB Technology's deployment measurements indicate reductions in inference token consumption of approximately 40–50 per cent in relevant use cases because the model is provided with pre-structured, policy-anchored context rather than repeatedly reconstructing that context from dispersed material.

These are KB Technology's internal deployment measurements and should be treated as implementation evidence rather than an independently validated industry benchmark.

The saving is reflected differently under each deployment model: as a smaller metered bill for a hosted model or as fewer GPU-hours and less energy consumed by a self-hosted model. The policy-relevant point is that both represent a reduction in computational demand. The standards should therefore consider whether structured operating context can contribute to measurable compute and energy efficiency irrespective of the model deployment model.

This is complementary to the physical efficiency measures proposed by the consultation. Measures such as power usage effectiveness address efficiency within the data-centre environment. They do not establish how efficiently the resulting compute is being used by AI workloads. A national approach to AI infrastructure that regulates the supply of compute and energy without considering the *efficiency of demand* addresses only part of the equation. Structured operating logic capture is one potential demand-side lever available at the application and organisational layer.

The submission therefore recommends that the eventual standards recognise measurable reductions in compute, inference and energy demand resulting from structured operating context as a legitimate contribution to the consultation's energy and no-net-cost objectives.

Table 1: Token and compute efficiency indicators		
Ser	**Indicator**	**Figure**
1.	Token overhead in unstructured AI deployments (industry audit benchmark)	40–60% of inference tokens
2.	Token reduction with captured operating logic (KB Technology measured deployments)	40–50% reduction
3.	Average annual AI inference cost per major enterprise deployment (Gartner, 2025)	~USD $1.03m

Open weight vs frontier

As mentioned previously, most business-as-usual activity is well served by smaller, open-weight models operating inside a structured operating context, because a modest model given the organisation's actual decision logic can be more useful for a defined business process than a larger model given none.

The resource gap is wider than tokens alone. Industry benchmarks place unstructured, unclassified data, commonly termed *dark data* at around 54 percent of an average organisation's stored data holdings,[3] meaning more than half of what an undocumented frontier model must search to answer an enterprise query carries no governance, structure or authoritative status. An undocumented frontier model operating against that environment must expend materially more compute per useful answer than a smaller, open weight model operating against a documented, structured operating logic layer because on every query the undocumented model repeats the work of discovery and inference that a captured operating logic layer has already done once, permanently.

Open weight models paired with captured operating logic therefore require significantly less computational resource and by direct extension significantly less electricity and cooling water, than undocumented frontier models performing the same class of everyday government or enterprise task. This is a second demand side lever, additional to the token efficiency finding above and it should inform how the standards treat compute and resource efficiency as a design criterion for AI deployment, not only as a data centre cost to be offset.

This is a direct contribution to the paper's stated objectives.

Every wasted inference token is a wasted unit of the same electricity draw the renewable electricity obligation and no net cost principle are designed to manage. A national approach to AI infrastructure that regulates the supply of compute and energy without addressing the efficiency of demand for it is addressing half the equation. Structured operating logic capture is a demand side lever available now, at no call on the energy system the paper is trying to protect and this submission recommends it be recognised as such in the eventual standards.

3. Transparency, repeatability and operable safety disclosure

The paper's section on the security and safety benefits of AI training proposes that frontier labs granted authorisation to train in Australia disclose defined reportable incidents to relevant Australian authorities and asks what information AI developers should be required to share and how.

As currently framed this obligation addresses the model. It tells government what went wrong inside a frontier system. It does not tell government what that system was doing inside an organisation's actual decision-making at the time because most government and enterprise deployments of AI currently have no structured, auditable record of the operating logic and decision authority the model was operating against when the incident occurred.

Cyber security and AI governance failures in large organisations typically do not occur because a framework or policy is missing. They occur because the control exists on paper and is not enforced at the point of execution so there is nothing for an incident report to actually trace against.

A review of Commonwealth AI transparency statements conducted by Kitbag Consulting in September 2026 found that, none whether inference was processed onshore, whether the Commonwealth controlled the underlying model weights or the jurisdictional exposure of the supplier.[4]

The finding is significant because the proposed disclosure regime for frontier developers will not by itself close this gap. The missing information sits primarily on the deploying entity's side of the relationship.

Kitbag Consulting's parallel submission to the Department of Home Affairs, made in July 2026 on proposed amendments to the Security of Critical Infrastructure Act 2018,[5] addresses this same problem from the incident classification side and its central finding applies directly here: an entity cannot secure what it has not modelled. The cyber security exposure of an AI system is not a property of the model. It is a property of the model's position in the operation, meaning what triggers it, what information flows into it, who receives its outputs and which decisions those outputs affect.

That submission recommends AI systems be documented the way an operating model already documents a person, as a defined delegation with a named accountable human and recommends a system be treated as materially uncontrolled where it exceeds that documented delegation or where its specified human interruption mechanism fails.

This submission adopts the same logic for AI training disclosure. A structured operating logic layer is what allows an entity to say, credibly, what a reported incident actually did and what it was authorised to do, rather than reconstructing that picture after the fact from vendor-held logs the entity does not control.

The problem compounds as AI moves from single models to interacting agents. The Australian AI Safety Institute's first substantive report, Risks and Controls for Multi-Agent Systems, commissioned from the Gradient Institute and published in August 2026, finds that a collection of safe agents does not guarantee a safe collection of agents, because interactions between agents produce emergent behaviours and failure modes that no single agent's own testing reaches.

The report frames this around three governance tiers: singular governance, where one organisation governs every agent in an interaction; federated governance, where multiple organisations share an environment under agreed rules; and open environments, where no common governance binds the interacting agents at all.

An entity cannot establish which tier it is actually operating in or trace an incident to its cause, unless it can first state what its own agents are authorised to do and to whom that authorisation traces. That is precisely the structured, auditable record this submission is recommending and it becomes more necessary, not less as frontier AI training moves from isolated deployments toward the cross-organisational agent interactions the Institute's report identifies as the harder and less visible risk.

A captured, machine readable operating logic layer resolves this by design rather than by audit. Every action taken within it, human or AI is traceable to a delegated authority, a defined process step and a governed data source recorded as the action occurs rather than reconstructed afterward. Output produced under this structure is repeatable and consistent for a given input, because it is drawn from the same authoritative process and policy content every time, rather than reconstructed by inference on each query. This is the practical meaning of the Deming principle that quality should be built in, not inspected in, applied to AI governance.

This gap is measurable and already producing failure at scale.

MIT and RAND reported in 2025 that 85 percent of AI projects fail to meet their expected outcomes and a separate MIT study the same year found 95 percent of proof-of-concept AI programmes failed to deliver measurable financial return.[6] S&P Global found the share of enterprises abandoning most of their AI projects rose from 17 percent in 2024 to 42 percent in 2025.[7] Deloitte's 2026 research found only one in five companies has a mature governance model for autonomous AI agents, even as agentic AI usage accelerates.[8]

McKinsey has separately found 70 percent of broader transformation initiatives fail to achieve their stated objectives.[9] These figures describe the same structural gap from different angles. Organisations are deploying AI against operating environments whose decision rights, processes and authorities have never been made explicit, so the AI inherits the ambiguity rather than resolving it. Three prior generations of enterprise governance discipline, enterprise architecture, business process management and RACI style accountability models, each addressed part of this chain, from mission down to task ownership, but none completed it end to end into a form an AI system can directly execute. That gap, not a shortage of published ethical principles is what current responsible AI frameworks largely leave unaddressed.

This submission recommends that the disclosure obligations proposed under the security and safety conditions be extended, for large scale deployments of frontier AI within government and Security of Critical Infrastructure Act entities, to include a requirement that the deploying entity be able to demonstrate a structured, auditable record of the operating logic and decision authority governing the AI's use, including where its agents interact with agents outside the entity's own governance. This gives an incident report something to trace against and gives government visibility not only of what a model did but of what it was authorised to do.

4. The commoditisation and replacement risk and third party access

Frontier AI providers face substantial pressure to convert very high levels of infrastructure and research expenditure into durable commercial value.

Commoditisation

One potential consequence is commoditisation: the possibility that an organisation's operating model, accumulated decision practices and ways of executing work become progressively embedded within a vendor platform and are ultimately supplied back to the organisation as part of a proprietary subscription.

This is already visible in the integration of frontier models into enterprise software platforms. Salesforce has integrated frontier models into Agentforce. MYOB has integrated Claude and ChatGPT directly into its platform.

Where an organisation's ERP, CRM or ITSM platform embeds a foreign frontier model directly against its operating data, the provider may sit materially closer to the organisation's operating and decision context, rather than simply providing a discrete AI service. This presents a *potential* sovereignty and dependency issue because the architectural boundary between the organisation's own operating logic and the vendor's platform becomes less clear.

Replacement

The second is replacement. Compute at a scale no domestic operator can fund is used to reverse engineer and replicate an entire industry function, delivered faster and cheaper and eventually without the incumbent industry at all.

Australia has a recent, well documented precedent for the jurisdictional dimension of this risk. In December 2025, following formal risk assessments conducted after at least nine prior rejections across various federal agencies, the Swiss Armed Forces concluded that Palantir's software could not guarantee that sensitive military and government data and the operational logic built on it, would remain beyond the reach of United States intelligence agencies under US law.[10][11] The conclusion was not about Palantir's technical capability. It was that dependency on a single foreign platform for mission critical operating logic was an inherent feature of the product, not a manageable risk, regardless of where the servers sat.

The same jurisdictional exposure identified in that case, arising from the US CLOUD Act,[12] and its equivalent under China's National Intelligence Law and Data Security Law,[13] applies to any Australian entity running its operating logic inside a foreign owned platform, irrespective of where that entity's servers or data centres are physically located.

This is the practical meaning of third party access to enterprise and national ways of working that the paper's current disclosure framework does not capture. Data residency and model level incident reporting do not address it, because the exposure exists at the level of who defines and controls the logic by which decisions get made, not at the level of where the bytes are stored.

Kitbag Consulting's submission to the Department of Home Affairs on the Security of Critical Infrastructure Act reaches the same conclusion from a different direction, recommending that a third party holding the sole authoritative record of a function's operating logic, its processes, delegations and decision rules, be recognised as exercising material practical control over that function whether or not it holds any ownership interest in it. An operator that holds the only account of how an entity works holds the function itself, whatever the contract says about ownership.

This submission recommends the AI standards explicitly recognise commoditisation of operating logic, meaning the embedding of a foreign frontier model directly against an entity's operating and decision data, as a distinct risk category from a safety incident and require government and critical infrastructure entities to disclose where this embedding exists.

5. Aggregation, sovereignty and the definition the standards should adopt

The paper correctly frames hosting frontier AI training in Australia as a way to deepen the country's capacity to understand, influence, manage and respond to emerging risk and to give Australia a seat at the table in global AI safety and security standard setting. This logic applies with equal force to the operating logic layer and arguably with greater force, because operating logic is where actual decisions, not just model behaviour are made.

An organisation that does not own its operating logic exposes not just its own operations but in aggregate, the country's decision-making to a party that can observe how it decides, before it decides. Where enough government and enterprise operating logic runs on infrastructure subject to foreign jurisdiction, that visibility aggregates from an organisational risk into a national one, because the entities holding legal reach over that infrastructure gain a persistent, structural view into how Australian institutions actually function, independent of what data residency or model hosting arrangements are in place.

This submission recommends the eventual legislation define sovereignty, for the purposes of the AI standards, as control of operating logic, meaning the decision rights, policies, processes and institutional knowledge governing an organisation's actions, rather than by data residency or the jurisdiction in which a frontier model is hosted or trained. Data residency is a necessary condition. It is not a sufficient one.

6. What operating logic capture requires in practice

Operating logic capture does not require replacing existing enterprise systems. Systems of Record, meaning ERP, CRM and ITSM platforms, remain the authoritative source for transactional data and should stay exactly where they are.

What is required is a separate, Australian owned governance layer that references and orchestrates those systems through controlled, task scoped associations rather than deep, privileged integration, so that operating logic, meaning the mission, functions, decision authorities, processes and policy linkages that define how the organisation works, is held independently and does not disappear when a vendor changes terms, incorporates foreign AI, exits the market or is acquired by a foreign entity.

The cost of not doing this is already a matter of public record in Australia. The whole of government GovERP program, which sought to standardise corporate systems across roughly one hundred agencies on a single enterprise platform, was wound back after expenditure the Minister for Finance described as approximately $390 million and was decommissioned in January 2025.[14] The Queensland Health payroll replacement required more than 2,500 system customisations to encode the department's own pay rules, affected the pay of tens of thousands of staff and was estimated by the subsequent Commission of Inquiry to cost in the order of $1.2 billion to rectify and operate over eight years.[15]

The Department of Defence is presently consolidating approximately 500 applications onto a single enterprise resource planning platform at a cost valued in the billions, a program still in delivery.[16] These programs differ in their particulars. Together they show one pattern: once an organisation's operating logic is configured directly inside a large execution platform rather than held separately from it, substitution becomes slow, expensive and difficult to reverse and the cost is paid regardless of which vendor or platform was chosen.

The sequence matters as much as the architecture. Where a sovereign operating logic layer exists before a major ERP or AI investment is made, that investment draws from and references the agency's own logic rather than becoming its only record. Where the layer is absent, each major platform reconstructs the organisation's operating logic inside itself, one procurement at a time and command of that logic is ceded by default.

KB Technology's DOLIUM is one existing example of this architecture, Australian owned, IRAP assessed and listed on the Digital Transformation Agency's Software Marketplace Panel. It is cited here as evidence that the approach is deliverable today, not as the only implementation the standards should recognise.

Conclusion

This submission has argued that the operating logic layer, not the physical infrastructure or the frontier training conditions the paper currently regulates, is what determines whether Australia captures the economic and governance benefit of its own AI investment. The physical layer secures the energy, water and workforce this build-out requires. The behavioural layer secures safety and security commitments from the labs training on Australian soil. Neither layer, on its own, secures the thing that actually determines who benefits: control over how government and enterprise decide and act.

The six findings set out above are one argument, not six separate ones. Frontier access is not required for the substantial majority of everyday AI use, so Australia's dividend does not depend on winning the frontier race, and it does not depend on frontier training landing onshore either, because the continuity risk this paper is trying to manage through location is a jurisdictional risk that location alone cannot fix. Captured operating logic is what delivers the token, compute and energy savings that reinforce the standards' own no net cost objective. It is what turns the paper's disclosure obligations from a model-level report into something an incident can actually be traced against.

Left uncaptured, it is what is already being commoditised into vendor subscriptions and exposed to undisclosed third party access, the same exposure Australia has already seen materialise in a foreign jurisdiction's treatment of a comparable platform. The cost of getting the sequence wrong, capturing this layer after a major investment rather than before it is not hypothetical. It is already a line item in the Commonwealth's own accounts.

None of this argues against the paper's core program. Large data centre investment, frontier training conditions and the local capability mechanisms the paper proposes are all necessary, and this submission supports them. What it argues is that the standards, as currently scoped, regulate the infrastructure AI runs on and the models that run on it, without regulating the layer that determines whether the country running that infrastructure actually controls the decisions being made on top of it. That gap is closeable now at the design stage before the same pattern that produced GovERP, the Queensland Health payroll rebuild and the Defence ERP program repeats itself inside an AI standard instead of an enterprise system.

The recommendations above set out seven specific ways the standards can close it. Adopted together, they would make Australia's AI infrastructure investment do double duty, building the physical and behavioural conditions the paper already proposes, while also building the one layer that determines whether the economic and governance benefit of that investment stays in Australia or moves offshore with the next change in a foreign provider's terms.

Recommendations

This submission recommends the AI standards:

1. Define sovereignty, throughout the eventual legislation, as control of operating logic to mean an organisation's decision rights, policies, processes and institutional knowledge, rather than by data residency or model hosting location alone.
2. Recognise that the substantial majority of government and enterprise AI use does not require frontier models and ensure the standards do not effectively position frontier training access as a precondition for Australia's AI productivity dividend.
3. Extend the proposed AI training security and safety disclosure obligations to require large scale deployments of frontier AI within government and Security of Critical Infrastructure Act entities to demonstrate a structured, auditable operating logic layer against which incident reports can be traced.
4. Recognise the embedding of a foreign frontier model directly against an entity's operating and decision data, through Enterprise Resource Planning, Customer Relationship Management or IT Service Management platforms, as a distinct disclosure category from a model safety incident and require government and critical infrastructure entities to report where this embedding exists.
5. Recognise measurable inference token and energy efficiency, delivered through structured operating context, as a legitimate contribution under the local capability and research investment mechanisms currently being considered, given its direct bearing on the energy and no net cost objectives set out in Part 2 of the paper.
6. Require entities deploying agentic or autonomous AI at scale within government to demonstrate governance maturity against a defined operating logic standard, rather than relying on adherence to general AI ethics principles alone, given independent research finding only a minority of organisations currently hold a mature governance model for autonomous AI agents.
7. Align the sovereignty definition and assessment approach in these standards with the tiered sovereignty framework already proposed to the Digital Transformation Agency for coordinated contracting arrangements,[17] so that a single, consistent sovereignty methodology, covering jurisdictional exposure, operating logic ownership and Australian economic contribution, applies across Commonwealth AI procurement rather than being assessed differently instrument by instrument.

[1] Kitbag Consulting, submission to the Joint Select Committee on Artificial Intelligence on the Commonwealth's jurisdictional exposure in AI procurement, 2026.

[2] Axios, "Trump administration blocks foreign access to Anthropic's most powerful AI" (12 June 2026); ABC News, "Trump administration forces Anthropic to pull models worldwide" (13 June 2026); access restored 1 July 2026.

[3] KB Technology Pty Ltd, DOLIUM — Enterprise System of Work: Business Case for Defence Data Division, 2026 (dark data benchmark).

[4] Kitbag Consulting, submission to the Joint Select Committee on Artificial Intelligence, September 2026 (review of 73 Commonwealth AI transparency statements).

[5] Kitbag Consulting, submission to the Department of Home Affairs on Measures 5 and 17, Security of Critical Infrastructure Act 2018 Tranche 2 consultation, 31 July 2026.

[6] MIT and RAND, AI project outcome research, 2025, as cited in KB Technology Pty Ltd, DOLIUM Powerplay: AI Graveyard or AI Nirvana, March 2026.

[7] S&P Global, enterprise AI project abandonment research, 2025, as cited in KB Technology Pty Ltd, DOLIUM Powerplay: AI Graveyard or AI Nirvana, March 2026.

[8] Deloitte, autonomous AI agent governance maturity research, 2026, as cited in KB Technology Pty Ltd, DOLIUM Powerplay: AI Graveyard or AI Nirvana, March 2026.

[9] McKinsey, transformation initiative outcome research, as cited in KB Technology Pty Ltd, DOLIUM Powerplay: AI Graveyard or AI Nirvana, March 2026.

[10] Republik (Switzerland), investigation into Swiss federal agencies' rejection of Palantir Technologies, based on 59 freedom of information requests, reported December 2025.

[11] TechDirt, reporting on the Swiss Armed Forces' internal risk assessment of Palantir Technologies, December 2025.

[12] Clarifying Lawful Overseas Use of Data Act (US CLOUD Act), 2018.

[13] People's Republic of China, National Intelligence Law, 2017 and Data Security Law, 2021.

[14] Digital Transformation Agency, GovERP Reuse Assessment, August 2024; Australian Government Minister for Finance, public statement on GovERP expenditure and decommissioning, 2025.

[15] Queensland Health Payroll System Commission of Inquiry, Final Report, The Honourable Richard N Chesterman AO RFD QC, 31 July 2013.

[16] Department of Defence, Defence Enterprise Resource Planning Program (SAP S/4HANA consolidation of approximately 500 legacy applications), in delivery.

[17] Kitbag Consulting, submission to the Digital Transformation Agency on the Coordinated Contract Agreements discussion paper, June 2026 (sovereignty tier framework and Deed of Sovereignty).

DOLIUM

What is Operating Logic?

Operating Logic is the authoritative digital representation of how an organisation operates: the decision rules, execution sequences, governance constraints and resource dependencies that translate strategic intent into operational action.

Own Your Operating Logic



With DOLIUM, Your Operating Logic Delivers:

Governance
Every decision, action and data flow is visible, auditable and accountable, giving you assurance that your organisation operates within the boundaries it has set.

Sovereignty
Your data, your processes, your AI: under your control. Sovereignty by design and independently verified, so you can operate with confidence at every level.

Operational Foundation
Your decision rules, authority structures and governance boundaries captured as sovereign IP, giving every AI use case the real context it needs to succeed and scale beyond a pilot.

Connection
Every system you already have: connected through open standards. Plug and play for ITSM, ERP, CRM to exploit the latest technology.

What is DOLIUM?

DOLIUM is an integrated digital environment that seamlessly links systems, data, documents and policy into a single intuitive interface to the data repositories, enabling AI-assisted decision making.

DOLIUM enables your operating logic, capturing how your organisation works and turning it into a foundation you own.

iAWARDS25 ACT WINNER EXCELLENCE IN INNOVATION

irap

AUSTRALIAN MADE AND OWNED

SCAN ME

GOVERN AI IN ALL ITS FORMS

DOLIUM

AI comes in various forms. AI agents, automated activities, generative AI and large language models are already operating across your organisation. Each is software that makes or supports a decision, and each depends on context to do it well.

Yet they often operate without context that explains what they do, what data they rely on, how they interact or who approves them.







Why Choose DOLIUM?

Operational View
- This is why we built DOLIUM. It provides a complete operational view of your AI ecosystem.
- Every agent, decision and data flow is visible, auditable and accountable, giving you assurance that AI operates within the boundaries your organisation has set.

AI Assisted Decision Making at the Speed of Thought
- Small language models delivering insights contextualised to your operations
- Air-gapped LLM deployment that keeps your IP inside your organisation
- Decisions informed, traceable and aligned to mission at the point of action

Responsible AI at Scale
- Orchestration of governable and transparent agentic AI
- Hyperautomation across the systems you already run
- Deploy hundreds to thousands of contextualised agents, with automation that is visible, auditable and repeatable

Use Case Identification and Generation
- AI readiness scoring across your operations
- Suggested improvements to lift operations to AI readiness
- Agents built from your evolving operating logic and business requirements
- AI use cases generated directly from your contextualised operations

DOLIUM AI Readiness Measures

Activity Maturity
How well the work is documented, standardised, and repeatable, the raw material any AI needs to learn from.

Controls Maturity
Whether the right policies, controls, and compliance are in place and genuinely connected to the work on the ground.

Accoundabilities
Whether roles and accountability are clear, who performs the work, who owns it, who signs off.

Measure your organisation's readiness for AI across process, governance and people.

DOLIUM

Email Us
dolium@kbtechnology.com.au

Visit Our Website
dolium.ai

AUSTRALIAN MADE AND OWNED