



A thought leadership series
from NCS Australia:

Perspectives on AI & technology for government





Four opinion papers addressing the real challenges facing government technology leaders in 2026.

Paper 1 - AI in government: the procurement problem nobody's talking about

Paper 2 - Sovereign capability is not a data centre

Paper 3 - Workforce is the strategy

Paper 4 - Platform, people, and process: why ai governance is the missing piece

AI in government: The procurement problem

There is no shortage of enthusiasm for artificial intelligence in the Australian public sector. Across every portfolio, agencies are exploring use cases, standing up pilots, appointment of Chief AI Officer (CAIO) roles and fielding briefings from vendors who promise transformational outcomes. The potential is exciting. The energy is real. So is the problem.

The problem is not technology. It is not policy, which is catching up. The problem is procurement and the time it takes. By the time a standing offer panel is

established, the technology listed on it is already a generation behind.

When public sector procurement moves in years but artificial intelligence advances in weeks, standing offer panels don't mitigate risk - they institutionalise obsolescence.

Australia's government procurement frameworks were designed for a world in which technology was stable, vendors were distinct, and evaluating capability was a matter of comparing specifications, implementation timeframes and delivery risk. AI invalidates all three assumptions simultaneously.

The pace of development in AI is unlike anything the enterprise technology market has previously experienced. A foundation model that represented the frontier six months ago may be superseded today and not incrementally, but potentially fundamentally. Procurement cycles that take 12 – 18 months to contract are, by definition, evaluating yesterday's technology. Agencies that lock into

arrangements may find themselves contractually tethered to a solution or vendor whose AI capability has been leapfrogged before the first purchase order is raised.

The second problem is vendor distinction. Traditional procurement assumes you can identify and evaluate discrete suppliers. AI has made that increasingly difficult. The large language model underlying an agency's chosen product may be licensed from a third party. The infrastructure may sit on a hyperscaler. The fine-tuning may have been done by a boutique specialist. The integration may be the only genuinely differentiating layer. Buying 'an AI solution' from a single vendor is often a false construct but procurement frameworks still treat it as real.

The third assumption surrounds capability evaluation. This breaks down entirely when the technology changes faster than evaluation cycles. By the time an assessment panel has assessed a product, scored it, challenged the scores, gone through probity review, and awarded the contract, the product has been updated three times and the market has moved on.

A different model is needed

The agencies making the most progress on AI are not the ones with the most sophisticated vendor engagement. They are the ones who have invested in understanding what they are trying to achieve, separated that question from the question of which technology delivers it, and have built strategic relationships that can support them in navigating a market that will keep changing.

That means outcome-based procurement beginning with defining what success looks like. Not assessing products. It means procurement structures that allow flexibility across vendors and generations of technology. It also means engaging vendors who have no stake in which AI platform wins, only in whether the agency's objectives are met.

NCS works with government agencies on exactly this basis. We are not the vendor of an AI product. We are the partner who helps agencies define the problem, evaluate the options and the roadmap, build the capability to adapt as those options change tomorrow, and integrate whatever solution is chosen into the broader landscape the agency operates in. The question is not 'which AI solution should we buy?' It's 'what are we trying to achieve, and how do we stay capable of achieving it as the technology evolves?'

The imperative is not to acquire a fixed AI product, but to anchor technology to a clear institutional purpose, and build the systemic resilience required to sustain it across generations of models.

This is a harder question. It requires more than a briefing from a vendor. But it is the right question. Agencies that are asking it are the ones that will have something to show for their AI investment in three years' time.

Sovereign capability Is not a data centre

Data sovereignty has become one of the most discussed concepts in Australian government technology circles. Recent years have surfaced geopolitical tension, critical infrastructure vulnerabilities, and a sharper national focus on strategic dependencies. Because of this, the instinct to keep data onshore is understandable. In many contexts, it is also correct.

But there is a conflation happening that, if left unchallenged, will leave agencies significantly exposed. Governments are treating data residency as if it were the same thing as capability sovereignty.

It isn't, and the difference matters enormously.

Keeping data onshore means nothing if the skills, architecture decisions, and integration knowledge all live offshore.

True digital resilience cannot be outsourced. If your data is onshore but your architectural memory and integration capabilities live offshore, you have retained the liability while exporting the capability.

Data residency is a question of where data is stored and processed. It is a necessary condition for sovereignty in certain classification contexts, but it is not sufficient. A government agency can store every byte of data on Australian soil, in an Australian-certified facility, and still be profoundly dependent on a foreign hyperscaler's engineers to understand how it is structured, a global systems integrator's offshore team to maintain the integrations that connect it, and a vendor's product roadmap for a global market to determine what can be done with it.

That is not sovereignty, it is residency. And the distinction matters because residency is easy to achieve and easy to audit, while genuine capability sovereignty is hard to build and harder to measure and maintain. The risk is that agencies tick the residency box and believe they have addressed the strategic problem when they have only addressed the compliance problem.

What genuine capability sovereignty looks like

True sovereign capability has three components that are frequently absent from the current conversation.

The first is human capability. The skills, knowledge, and judgment that exist within Australian institutions and with Australian-based partners. This includes both the agency's own workforce and the ecosystem of providers it relies on. If every critical architecture decision requires escalation to an offshore centre of excellence, sovereignty is illusory regardless of where the data sits.

The second is architectural independence. The ability to understand, modify, and if necessary, migrate or replace the technology that underpins critical functions. Agencies that have allowed deep, undocumented dependencies to accumulate solely with vendors have, in practice, ceded architectural sovereignty regardless of their contractual rights. You can own the data and still be unable to move it.

The third is continuity of knowledge. The institutional understanding of why systems were designed the way they were, what assumptions were made, and what would need to change if the operating environment shifted. This is perhaps the most fragile form of sovereignty, because it exists in people, not in documentation, and it walks out the door when key individuals leave or when engagement with a vendor ends. You can own the data and still be unable to move it. That is not sovereignty, it is a more expensive form of dependency.

Data ownership without portability is an expensive illusion. If your architecture keeps information immobile, you haven't realised sovereignty.

NCS approaches sovereign capability as a deliberate design objective, not a compliance checkbox. As a company headquartered and operationally grounded in the Asia-Pacific region, with security-cleared Australian practitioners embedded in long-term government engagements, we build capability transfer into the way we work. The goal is not to make clients more dependent on NCS. It is to leave them with more capability than they started with. A capability that is understood, documented, and actively transferred.

The conversation about sovereignty in Australian government is important and necessary. But it needs to move beyond the data centre. The question is not just where your data lives. It is whether you have the knowledge, the skills, and the architectural understanding to independently and sustainably do something with it on your own terms.



Workforce is the strategy

Ask any government technology leader what their biggest challenge is, and the answer is rarely the technology. It is people. Skills shortages, capability gaps, difficulty attracting and retaining talent in a market where the private sector offers more flexibility and often more money. Every agency feels it. Almost every programme plan treats it as a constraint to be managed rather than a strategic problem to be solved.

That framing is costing government significant amounts of money and producing technology programmes that deliver less than they should.

Government technology programmes are persistently designed around the technology, with workforce considerations bolted on at the end as a change management workstream

Designing public sector tech programmes around platforms rather than people treats workforce capability as a trailing administrative afterthought.

The standard model for a major government technology programme goes something like this. A policy change or problem is identified and assessed. A business case is developed around the solution options. A procurement is run. A vendor or integrator is selected and engaged. Implementation begins. Somewhere in the middle of delivery, a 'change management workstream' appears. It runs training sessions, produces communications, and manages the transition. Then the programme closes, the integrator leaves, and the agency is left with a system it understands incompletely and a workforce that is operating it at partial capacity.

This is not a description of failure. It is a description of normal.

Flipping the model

The agencies that get the most sustained value from technology investment are not necessarily the ones that chose the best technology. They are the ones that understood before the procurement, what their workforce needed to be capable of doing, and designed the programme to achieve that outcome alongside the technology outcome.

This means starting with capability mapping, not solution selection. What do people in this agency need to be able to do that they cannot do today? What skills exist, where are the gaps, and which gaps are genuinely addressable through training versus which require a different resourcing model? The answers to these questions should shape the procurement, not follow from it.

It means building skills transfer into contracts as a deliverable, not an afterthought. If an integrator is engaged to implement a platform, the contract should specify (along with measurable outcomes) how much of the knowledge required to operate, maintain, and evolve that platform will reside within the agency at the end of the engagement. Too few contracts do this. Too few agencies hold their vendors to account for it when they do.

If an integrator is engaged to implement a platform, the contract should specify how much of the knowledge required to operate and evolve it will reside within the agency at the end. Too few contracts do this.

An integration contract that fails to explicitly mandate internal knowledge transfer isn't an investment in capability - it is a long-term subscription to vendor dependency.

And it means rethinking what 'delivery' means. A programme that delivers a working system but leaves the agency dependent on external support to operate it has delivered the technology, but not the capability. In a resource-constrained environment, that distinction has direct budget consequences every year thereafter.

Why this matters now

The pressure on government workforces is not going to ease. AI will automate some functions and create demand for new ones. Agencies that have invested in genuine workforce capability including the ability to understand, adapt, and evolve their own technology environment, will be positioned to take advantage of that shift. Agencies that have accumulated a series of systems they do not fully understand will find it harder and be continually reliant on the external workforce market.

NCS works with government agencies on programmes where workforce capability is treated as a first-class outcome, not a workstream. That means co-designing solutions with the people who will operate them, building training and knowledge transfer into delivery milestones, and measuring success not just by whether the technology went live, but by whether the agency can drive it forward independently.

The technology in any programme will change. The next AI model, the next platform version, the next architectural standard will all arrive. The agencies that will adapt well are the ones with workforces that understand the fundamentals and not just the current implementation. That capability does not appear at go-live. It must be built deliberately, from the beginning.

Platform, people, and process: Why AI governance Is the missing piece

Every conversation about AI in government eventually arrives at the same point. An agency has selected a platform or is close to doing so. It has a team or is assembling one. It has a mandate, a use case, and in many instances genuine executive commitment. And yet the programme stalls,

underdelivers, or produces results that cannot be repeated, audited, or scaled.

The explanation is rarely the platform. It is rarely the people either, though skills are always a factor. The consistent missing ingredient across agencies and programme types is the operating framework. The structured discipline that determines how the AI is used, how it is deployed, what is permitted and what is not, how decisions are made and recorded, and who is accountable when something goes wrong.

Without that framework, platform and people are necessary but insufficient. And in a government context, insufficient is not good enough. The platform is the engine. The people are the drivers. But without the road rules, you have a very fast vehicle with no reliable way to know where it's going.

Platforms offer velocity and people provide the drive. But without governance, high performance is simply a faster way to get lost.

Why platform alone is not the answer

The AI platform market has matured rapidly. Agencies now have access to capable, enterprise-grade tools, large language models, automation platforms, analytics engines which would have been unimaginable five years ago. The vendors behind these tools are sophisticated, their security postures have improved, and in many cases the technology itself is genuinely impressive.

But a platform is an enabling condition, not a solution. Choosing the right AI platform without a framework to deploy and govern its use is analogous to purchasing a fleet of vehicles without road rules, licensing requirements, or maintenance schedules. The vehicles work. The outcomes are unpredictable.

In a government context, unpredictability has consequences that the private sector can absorb more easily. An AI system that produces inconsistent outputs in a commercial setting is a product problem. An AI system that produces inconsistent outputs in a regulatory, welfare, or law enforcement context is a governance and policy failure. The tolerance for ungoverned AI in government is not just lower than in industry. It is, or should be, approaching zero.

Why resourcing alone is not the answer either

The instinct, when a technology programme underperforms, is to resource the problem. More people, more specialists, more budget. Sometimes this is the right diagnosis. More often, it treats a symptom while leaving the underlying condition unaddressed.

Additional resourcing into an ungoverned AI environment does not produce better outcomes. It produces more output with the same structural vulnerabilities. If the framework for how AI decisions are made, reviewed, and recorded does not exist, adding people accelerates the rate at which ungoverned decisions accumulate. That is not progress. That is risk amplification.

The agencies that have made the most sustainable progress on AI adoption are not uniformly the best-resourced. They are the ones that invested early in defining how AI would operate within their environment, what guardrails would apply, what human oversight would look like at each stage. Also how they would demonstrate to auditors, ministers, and the public that AI-assisted decisions were sound. More resourcing into an ungoverned AI environment does not produce better outcomes. It produces more output at greater speed, with the same structural vulnerabilities.

Injecting resources into an ungoverned AI environment does not optimise outcomes. It merely accelerates risk, producing flawed data at a faster velocity while leaving foundational vulnerabilities perfectly intact.

The framework as bedrock

A governance framework for AI does several things that neither platform nor people can do independently. It makes the rules explicit. Not implicit, not assumed, not dependent on the judgment of whichever practitioner happens to be working on a given task. It creates an audit trail: a record of what was decided, on what basis, by whom, and with what level of confidence. It defines the gates at which human oversight is mandatory, and the conditions under which AI-generated output can proceed without additional review. And it provides a consistent operating model that does not degrade when personnel change, when a vendor updates their product, or when a programme scales beyond its original scope.

In short, a framework converts AI from a tool that works when conditions are favourable into a capability that performs reliably under real operational conditions.

PACE: A framework built for this environment

NCS has developed PACE: Plan, Ask, Check, Execute. PACE is an enterprise-grade AI governance framework designed specifically for the complexities of real delivery environments. PACE is not a set of principles or a policy document. It is an operational methodology that installs governance directly into the way AI is used across a programme or agency.

The four-stage discipline is straightforward in its logic. No AI output - no code, no document, no recommendation - is produced without an explicit, approved plan that names scope, applicable standards, and the evidence required to declare completion. Before execution, sub-agents and reviewers interrogate the plan and surface open questions rather than bury them. Security, compliance, quality, and policy gates run against the plan before work proceeds. And when AI produces work, humans verify outputs against defined evidence requirements with reviews and sign-offs committed to a traceable audit trail.

P - Plan: Define scope, standards, risk, and accountability before anything is produced. No AI output without an approved plan.

A - Ask: Interrogate assumptions. Surface open questions. Apply expert review and confidence scoring before execution begins.

C - Check: Run security, compliance, quality, and policy gates. Register risks. Where a gate fails, work does not proceed.

E - Execute: AI produces work under explicit constraints. Humans verify against plan requirements. Decisions are recorded and traceable.

PACE is platform-agnostic and IDE-agnostic. It works with the tools an agency already has or is procuring, not as a replacement for them. It installs governance rules and workflows directly into a project environment, creating always-on guardrails that the AI

must operate within. It includes a risk framework - RIPM (Risk, Impact, Proximity, Mitigation) - that requires every plan to carry a risk register with mandatory human sign-off before execution proceeds. And it includes a confidence scoring mechanism that tags every AI output with a certainty level, triggering mandatory human review when that certainty falls below defined thresholds.

Critically, PACE is also designed to scale. The same four-step discipline applies whether the task is rewriting a function or designing a migration strategy. The inputs and outputs change; the governance structure does not. That consistency is what makes it suitable for government environments, where the operating scale, the regulatory context, and the public accountability obligations demand something more robust than good intentions and skilled practitioners. PACE converts AI from fast-but-risky into fast-and-dependable - across the full lifecycle, not just in favourable conditions

PACE bridges the gap between raw AI velocity and operational resilience - ensuring high-speed innovation remains dependable through every lifecycle stage, not just under ideal conditions.

Bringing it together

The argument across this series of papers has been consistent. AI procurement frameworks need to be redesigned for a technology that changes faster than panels can be established. Sovereign capability requires more than data residency, It requires human knowledge, architectural independence, and continuity. Workforce strategy needs to precede technology selection, not follow it.

Paper 4 is the synthesis. Platform, people, and process are not three separate investments to be made in sequence. They are three interdependent conditions that must be present simultaneously for AI adoption in government to produce outcomes that are reliable, repeatable, and defensible.

NCS brings all three. We are vendor-neutral on platform. Our role is to help agencies choose and integrate the right technology for their context, not to sell them a product. We are committed to building genuine workforce capability and not create a dependency. And we bring PACE as the operating framework that holds the model together: the governance discipline that makes AI fast and dependable rather than just fast.

For agencies that are serious about AI adoption - not as a pilot, not as a proof of concept, but as a sustained operational capability - this combination is what the next stage of the journey requires.

Scott Gledhill is Executive Director, Federal Government at NCS Australia. NCS is a leading technology services company headquartered in Asia-Pacific, with deep capability in AI, cloud, and digital transformation for government clients. NCS's PACE framework is an enterprise-grade AI governance methodology available to government agencies and delivery partners.



Scott Gledhill | Executive Director, Government Sector | **NCS Australia**

Scott is a driven executive leader with over 20 years of experience executing growth strategies and long-term account planning across global technology enterprises, including IBM, Accenture, and Salesforce. A focus in cross-sector portfolio management and complex pursuits, mentoring high-performing teams to optimise public sector strategies, operational processes, and client experiences.

For more information, visit www.ncs.co/en-au

