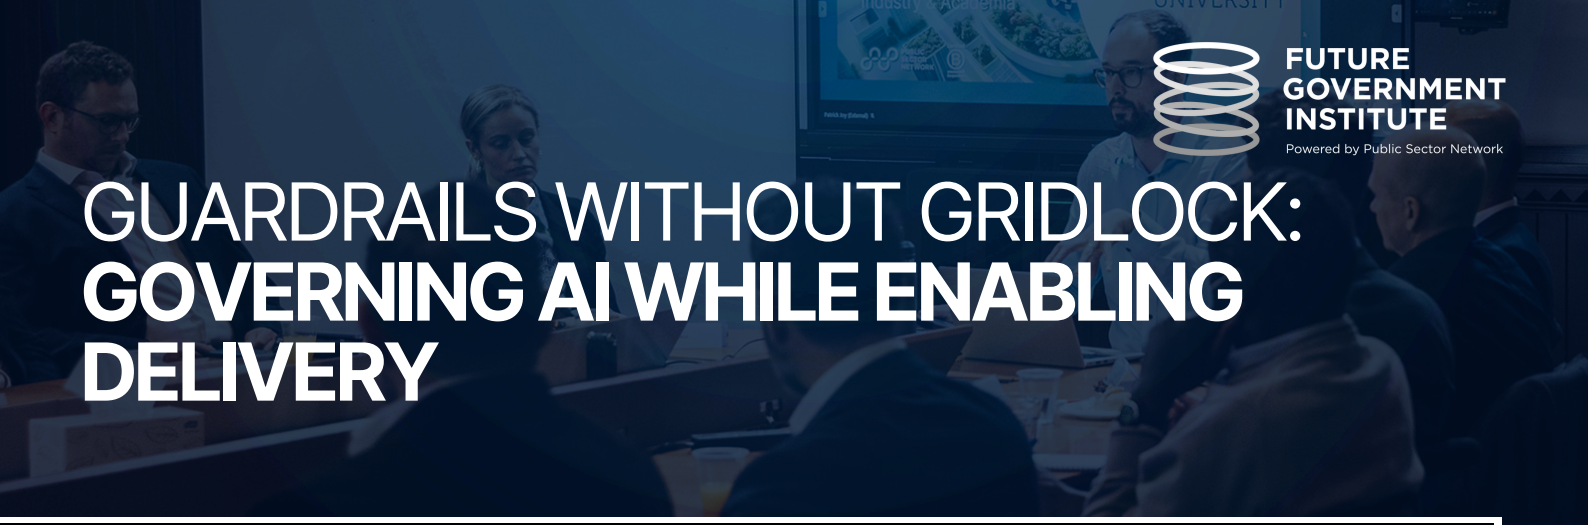

FUTURE GOVERNMENT INSTITUTE
Powered by Public Sector Network

GUARDRAILS WITHOUT GRIDLOCK: GOVERNING AI WHILE ENABLING DELIVERY

EXECUTIVE SUMMARY | VOL.36 | NEW SOUTH WALES

UTS UNIVERSITY OF TECHNOLOGY SYDNEY

IN BRIEF

The Research Council convened a roundtable focused on moving from AI ambition to delivery, with participants converging on a blunt reality: most agencies are not short on pilots, frameworks, or vendor options, but they are short on repeatable pathways that turn experimentation into sustained outcomes. Across the discussion, "rebuilding the plane while flying it" emerged as a shared metaphor for modernising legacy environments while maintaining essential services and risk tolerance in high visibility domains.

Participants described a shift from strategy to execution, emphasising product-led delivery, platform thinking, and pragmatic controls that reduce unsafe workarounds. A recurring insight was that "guardrails" are only effective when paired with usable, sanctioned tools, otherwise staff will route around restrictions via consumer AI. Several agencies are trialling tightly scoped deployments (for example, ringfenced copilots and controlled data environments), and using innovation labs or task forces to triage problems into prototypes within two-week cycles.

The roundtable also surfaced structural inhibitors: procurement cycles that outlast technology relevance, difficulty quantifying public value (particularly in health), and uneven executive understanding of AI's practical risks and benefits. The most actionable takeaway was to reframe AI as a productivity and service-improvement lever, grounded in baseline measurement and process redesign, rather than a standalone program.

By **Patrick Joy** | Head of Research & Analysis | **Public Sector Network**

KEY THEMES AND INSIGHTS

DELIVERY IS THE NEW DIFFERENTIATOR

Participants were clear that the sector has moved past debating whether AI matters, and is now wrestling with how to deliver without destabilising core services. Many agencies described a backlog of proof of concepts and pilots that generated learning but did not scale, often because they lacked a sustained operating model, funding continuity, or clear ownership. The shift in emphasis was towards delivery patterns that can be repeated: small, meaningful pilots; parallel run where risk is high; and disciplined scaling only after evidence is demonstrated.

A useful framing emerged: the goal is not to change government "all at once", but to improve outcomes within an agency or a bounded part of a system. This aligns with product thinking and capability embedding, where teams build confidence through incremental wins that can later be replicated.

"GUARDRAILS" WORK ONLY IF PEOPLE HAVE A SAFE DEFAULT

A strong, practical view emerged that restrictive policies alone do not stop shadow AI use. If staff can access a browser, they can upload documents and prompt external tools regardless of internal rules. The more effective approach discussed was behavioural and architectural: provide a safe, usable default (such as approved copilots or controlled internal environments), so the path of least resistance is also the safer path.

Examples included tightly ringfenced copilots, progressive expansion of access as confidence grows, and layered controls (identity, access, labelling, and tenancy protections). The intent is not to eliminate risk, but to reduce unmanaged exposure and create visibility over usage patterns.

SIMPLIFY THE TECHNOLOGY FOOTPRINT BEFORE YOU "AI-IT"

Participants repeatedly linked AI readiness to technology simplification. Where environments contain duplicated tools, overlapping platforms, and inconsistent standards, the cost of integration rises quickly and undermines any productivity gains AI might offer. Several comments echoed the lesson from early agile adoption: empowering teams without underlying standardisation can lead to a sprawl of unsupported tools and duplicated capability.

A practical example was the deliberate coexistence of two signature tools: one standard enterprise option for broad use, and one specialised tool for complex workflows. The point was not tool count, but conscious governance and clarity about why exceptions exist. In this framing, AI amplifies both strengths and weaknesses: simplified environments enable faster adoption and lower integration costs, while complex environments make every new model or agent more expensive to operationalise.

PLATFORM THINKING AS A ROUTE OUT OF PILOT SATURATION

Participants described emerging platform approaches, particularly for data and model access, as a way to move beyond isolated pilots. A common ambition was to establish enterprise data platforms where datasets are consolidated and governed, and where AI experimentation occurs inside controlled environments rather than through ad hoc external tools. This was framed as enabling reuse, reducing duplication, and creating a common foundation for retrieval, augmentation, and agent-like capability.

However, the group also noted that platform approaches require sustained leadership endorsement and multi-layer alignment. Where support exists from senior leadership, platform investment is achievable. Where it does not, pilots proliferate because they are easier to approve and easier to stop.

CHALLENGES AND BARRIERS

PILOT-TO-PRODUCTION BOTTLENECKS

Several participants described pilots that were “ready to implement” but stalled during assurance, endorsement, and governance pathways. This created fatigue and reinforced the perception that innovation is allowed only as experimentation, not as operational change. A related barrier was the absence of baseline measurement, making it difficult to prove impact and prioritise scaling.

QUANTIFYING VALUE IN COMPLEX DOMAINS

Health was raised as a case where value is not easily expressed in dollars. When benefits are framed as lives saved, reduced risk, or improved clinical outcomes, translating this into investment cases becomes harder, even when performance uplift is significant. Without accepted valuation models, solutions remain stuck as pilots.

PROCUREMENT CYCLES THAT OUTLAST RELEVANCE

Participants described innovation procurement pathways taking 8–12 months or longer, by which point chosen tools can become outdated. The deeper issue was translation: moving from a technical outcome to a business case and then into procurement rules requires people who can connect all three, and that capability is scarce.

LEADERSHIP CAPABILITY GAPS AND RISK BIAS

The group observed an uneven executive understanding of AI, often skewing towards risk stories and high-profile failures. This can lead to an overemphasis on restrictions or policy artefacts without corresponding investment in capability, training, and delivery pathways.

INTER-AGENCY SILOS AND LEGACY CONSTRAINTS

Even when processes are similar across agencies, access limitations, legacy systems not designed for integration, and differing incentives block shared solutions. The result is repeated reinvention and duplicated effort.

FUTURE FOCUS AREAS

PRACTICAL AI DELIVERY PLAYBOOKS FOR GOVERNMENT

A capability stream could codify repeatable delivery patterns that move initiatives from pilot to production, including baseline measurement, parallel-run design, operational readiness, and change management. This would help agencies avoid reinventing delivery mechanics for each new use case.

DATA PLATFORMS AND RETRIEVAL-READY ARCHITECTURES

Participants consistently tied AI outcomes to governed, reusable data environments. Future roundtables could focus on how to design enterprise data platforms, implement data product models, and build retrieval-ready architecture that supports safe RAG-style capability without data leakage.

EXECUTIVE LITERACY FOR AI GOVERNANCE AND VALUE

A recurring need was to elevate AI understanding beyond technical teams. A future focus area could be short, practical executive sessions on AI risk, assurance, and productivity value, framed in ministerial and Cabinet language.

PROCUREMENT REFORM FOR FAST-MOVING TECHNOLOGY

The Council could convene procurement leaders and digital delivery leaders to redesign pathways for purchasing AI-enabled tools and services, including modular contracting, outcome-based procurement, and mechanisms to prevent “procurement latency” from becoming the hidden blocker of innovation.

WORKFORCE TRANSITION AND TASK REDESIGN

Rather than framing AI as job removal, participants focused on redeploying capacity to higher-value work. A capability stream could support task redesign, human-in-the-loop operating models, and workforce skills development for roles such as prompt engineering, model governance, and AI product management.

INNOVATIVE IDEAS AND CASE STUDIES

1. PARALLEL-RUN DELIVERY IN HIGH VISIBILITY ENVIRONMENTS

Where failure carries significant public consequences, participants described running new approaches in parallel with established systems. This allows confidence-building without risking service continuity. The approach is scalable as a pattern, even if each implementation is domain-specific.

2. CONTROLLED AI ENVIRONMENTS AS A HARM-REDUCTION STRATEGY

Rather than attempting to ban external AI use, some teams are enabling sanctioned tools across the workforce, accepting that unmanaged usage will otherwise occur. By providing a secure default, agencies can reduce uncontrolled exposure, improve visibility, and progressively expand capability through maturity stages.

3. AI-ASSISTED CODE MODERNISATION WITH MEASURABLE PRODUCTIVITY GAINS

Participants described using AI to accelerate the translation of legacy code into modern languages, reducing effort from weeks to hours for specific components and creating a plausible pathway to modernisation at scale. The underlying insight was that AI can create material productivity uplift when paired with clear technical constraints and a defined pipeline.

4. AGENTIC TRIAGE FOR ICT SUPPORT AND OPERATIONS

A practical use case raised was applying agent-like automation to IT service management: correlating incidents, searching knowledge bases, routing tickets, and reducing mean time to resolution. This was framed as a high-value target because it reduces repetitive work and improves service reliability.

5. TIME-BOXED “INNOVATION LAB” MODELS TO SURFACE REAL PROBLEMS QUICKLY

Participants described two-week prototype cycles, where business units bring a problem, co-design a low-fidelity solution, and receive a working prototype quickly. This model reduces wasted effort on poorly defined problems and tests value early, while creating a visible pipeline for prioritisation.

STRATEGIC OUTCOMES AND RECOMMENDATIONS

IMMEDIATE ACTIONS

- **Establish a safe default AI toolset** with clear guidance on permissible use and data handling. Focus on usability and adoption, not only restriction, to reduce unmanaged shadow AI behaviour.
- **Baseline before pilots**: require each pilot to document the current process, time spent, error rates, and service impact measures, so outcomes can be credibly compared.
- **Create a two-speed delivery model**: fast experimentation lanes for low-risk prototypes, and a structured production lane with defined assurance steps, so teams know what "good" looks like to scale.
- **Run a "use case triage sprint"**: shortlist a small number of high-value, well-scoped problems that can deliver measurable outcomes before year-end, rather than expanding pilot volume.

MEDIUM-TERM GOALS

- **Stand up an enterprise data and AI experimentation environment** (even if minimal), where datasets, retrieval, and model access are governed. Treat this as foundational infrastructure.
- **Reduce tool sprawl deliberately**: conduct a rationalisation pass to identify duplicated platforms and overlapping capabilities, prioritising simplification that reduces integration cost and accelerates delivery.
- **Formalise an AI task force or community-of-practice model** that channels staff interest into structured working groups and prioritised pilots, with clear criteria for selection and scaling.
- **Procurement pathway tuning**: introduce modular purchasing approaches, outcome-based statements of work, and faster evaluation cycles to avoid buying tools that are outdated by contract signature.

LONG-TERM VISION

- **Shift from project funding to product sustainability**: ensure AI-enabled services are funded as ongoing capabilities, not short-term pilots, particularly where benefits require continuous model tuning and operational support.
- **Normalise task redesign and redeployment**: build operating models that explicitly use AI to reduce repetitive workload and redeploy capacity to higher-value work, avoiding productivity gains being absorbed into new layers of administration.
- **Build executive confidence through evidence**: develop a library of validated public sector patterns, including what assurance looks like, what productivity gains are realistic, and what governance is proportionate to risk.
- **Enable cross-agency reuse**: create mechanisms for agencies to surface common problems and reuse proven solutions, focusing on shared processes even when outcomes differ, to reduce duplication and accelerate sector-wide improvement.

ABOUT THE FUTURE GOVERNMENT INSTITUTE (FGI) RESEARCH COUNCIL

We've been able to engineer a new program antithetical to the classical red tape, administration, and risk-aversion that impedes innovation.

Public Sector Network has recently launched The Future Government Institute (FGI), a global hub for forward-thinking public sector leaders, innovators, and practitioners dedicated to shaping the next era of governance.

Our mission is to empower government professionals with the tools, insights, and networks needed to drive meaningful transformation - https://publicsectornetwork.com/future-government-institute/

Leveraging our extensive connections, we are uniting the sharpest minds from government, academia, and industry via monthly research-driven roundtables, hosted at esteemed national centres of research, courtesy of university partners across Australia and New Zealand.

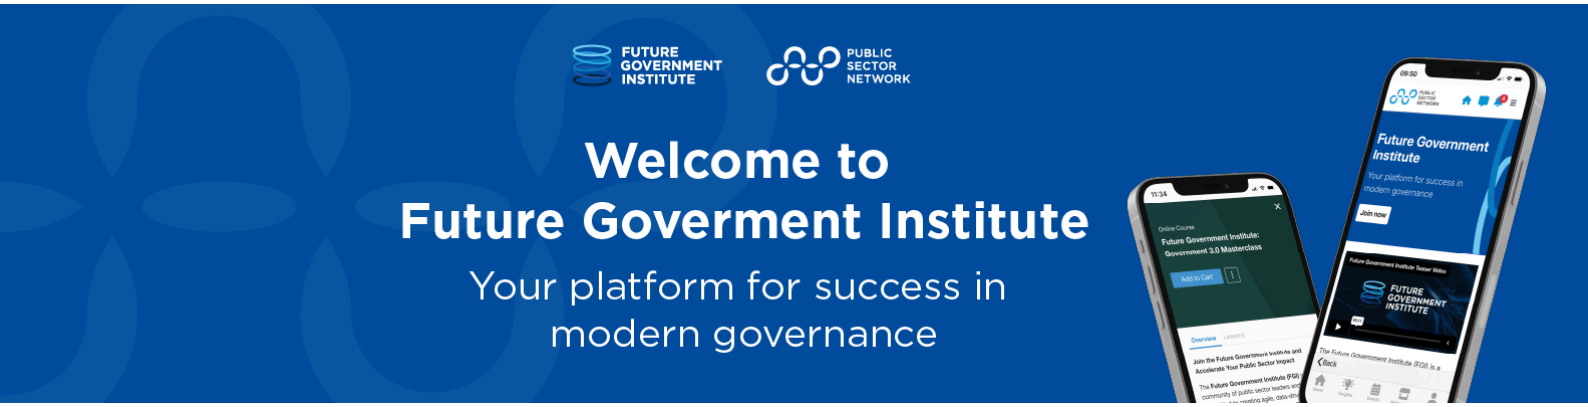

ABOUT PUBLIC SECTOR NETWORK

Public Sector Network is a research company that represents public sector professionals across Australia, Canada, New Zealand, and the USA. It develops roundtables, seminars, and conferences to suit current areas of interest to government agencies and their suppliers.

PSN's growing community spans across federal, state, and local government departments, healthcare, and education, allowing members to share information, access the latest in government innovation, and engage with other like-minded individuals on a secure and closed-door network.

AUSTRALIA / NEW ZEALAND
P +61 2 9057 9070
E INFO@PUBLICSECTORNETWORK.COM.AU

USA
P +1 (647) 969 4509
E HELLO@PUBLICSECTORNETWORK.COM

CANADA
P +1 (647) 459 8904
E CONTACT@PUBLICSECTORNETWORK.CO

Public Sector Network (Australia) Pty Ltd
ABN - 46 617 870 872 20-40
Meagher Street, Chippendale, Sydney NSW 2008, Australia