FUTURE GOVERNMENT INSTITUTE
Powered by Public Sector Network

GOVERNING THE LEAP: RESPONSIBLE AI, WORKFORCE DESIGN, AND THE NEXT PHASE OF READINESS

EXECUTIVE SUMMARY | VOL.32 | VICTORIA

RMIT UNIVERSITY

IN BRIEF

The Research Council convened a roundtable on AI readiness and emerging technologies, focusing on how public sector and research organisations are moving from early experimentation to structured adoption. Participants revisited an earlier maturity check, noting that where many agencies previously self-rated around “3 out of 10”, capability is now more uneven and polarised, ranging from areas still laying foundations through to teams operating at “9 or 10” where they are building, deploying, or critically evaluating AI.

A clear pattern emerged: most current adoption is framed through efficiency and back-office enablement, with limited appetite for high-stakes, citizen-facing automation while trust, governance, and assurance remain immature. Practical gains are already visible. One federal research agency reported 100% staff access to Microsoft Copilot 365 and early measurement showing an average productivity improvement of 26 minutes per staff member per day. Alongside these gains, participants raised deeper questions about responsible use, sustainability (including energy demand and “non-frivolous” use), vendor reliance, workforce impacts, and the risk of accelerating output at the expense of judgement and learning.

The discussion highlighted that AI readiness is less a technology upgrade and more a capability shift, requiring leadership clarity, workforce design, and shared evaluation approaches to translate productivity into better services and better work.

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

KEY THEMES AND INSIGHTS

RESPONSIBLE BY DESIGN, NOT BY POLICY HEADER

Participants consistently treated "responsible AI" as a practical operating discipline, not a compliance artefact. One federal research agency described an early decision to block public generative AI tools to create space for governance, then establishing an AI transparency statement and responsible use policy before scaling access. Their policy emphasis included human accountability for outputs, citation discipline, awareness of hallucinations and deepfakes, and clear guidance on handling sensitive information. Several participants stressed that policy alone is insufficient without training that matches real workflows, including "what not to do" scenarios, and role-specific guidance for researchers, corporate staff, and frontline teams.

A notable evolution was the inclusion of sustainability within "responsible use". Participants described growing staff concern about energy consumption and broader environmental impacts, with one organisation proposing explicit language to discourage frivolous use and to require demonstrable value. This surfaced an important gap: principles are easy to publish, but hard to operationalise without decision thresholds, measurement, and shared guidance on what "responsible" looks like in practice.

EFFICIENCY IS THE ENTRY POINT, BUT NOT THE END STATE

Across agencies, AI activity is largely anchored in efficiency: faster drafting, summarisation, knowledge search, sentiment analysis, and automation of repeatable tasks. Participants pointed to a common adoption curve where most staff sit in the "middle", seeing value but lacking skill or confidence to use AI safely and well. Several organisations are now rolling out foundational elements in parallel, such as policies, work instructions, training, and contained internal chat environments, with targeted pilots embedded in specific systems (for example, knowledge management and accreditation workflows).

However, the group also challenged "efficiency for efficiency's sake". Participants argued that productivity gains should be treated as a means to improve service quality and staff experience, not a blunt cost-cutting lever. In practice, this means translating time saved into better client outcomes, improved accessibility, or more capacity for complex, human-centred work.

THE HUMAN FACTOR: JUDGMENT, LEARNING, AND COGNITIVE DRIFT

A strong thread was concern that AI can accelerate output without improving understanding. Participants warned of "intellectual atrophy" if organisations reward speed and volume over depth and judgement, particularly when AI-generated content is circulated and re-summarised repeatedly. This was framed as both a workforce capability risk and a leadership risk: when change moves faster than organisational learning, leaders may rely on shortcuts, while staff lose opportunities to develop foundational skills.

Academic participants reinforced that limitations remain material, especially in language reasoning and domain nuance. The group noted that different AI techniques have different strengths, and that "AI" as a catch-all term obscures important distinctions. This supports a more disciplined approach: matching the tool to the job, and embedding human oversight where judgement is required.

CAPABILITY, SOVEREIGNTY, AND THE VENDOR GRAVITY PROBLEM

Participants raised growing unease that access to frontier models is shaped by a small number of global vendors, and that government and universities may be pulled into dependency by cloud-hosted services and proprietary platforms. This was framed as a strategic issue, not just procurement or security. Some argued that "sovereign AI" should be treated as a policy and capability agenda, spanning models, infrastructure, data governance, and localised context via techniques such as retrieval-augmented generation. The group also acknowledged trade-offs: sovereign aspirations must contend with practical realities of cost, performance, and rapidly evolving tooling.

CHALLENGES AND BARRIERS

GOVERNANCE MATURITY AND CONSISTENT TRAINING

Many organisations are still building baseline controls, including responsible use policies, work instructions, and role-based training. The gap is not simply "having a policy", but making it usable at the point of work, especially when staff are under time pressure and tools are easy to access.

TRUST AND ASSURANCE IN FAST-MOVING TOOLS

Leaders described being challenged publicly by staff who compare tools, models, and vendors, and expect explanations of why a chosen solution is safe and fit for purpose. Building confidence requires transparent assurance processes and clear communication about what checks have been done "under the hood".

SUSTAINABILITY WITHOUT OPERATIONAL THRESHOLDS

Participants recognised sustainability concerns, but also questioned how agencies define decision points. Without measurement, thresholds, or shared standards, sustainability risks becoming a vague principle rather than a decision-making input.

WORKFORCE DISRUPTION AND HOLLOWING OUT ENTRY PATHWAYS

The group raised a structural risk that early-career, "learning the ropes" roles may be automated first, creating a capability gap over time. This is not only a labour market issue, but a pipeline risk for government's future digital and policy capability.

PLANNING HORIZONS COLLAPSING

Technology leaders noted difficulty maintaining traditional 12-month roadmaps when major capability shifts can occur in a quarter. This creates governance tension: agencies need strategic direction while remaining flexible in delivery.

FUTURE FOCUS AREAS

PRACTICAL RESPONSIBLE AI THAT CAN BE AUDITED

Host a capability stream on turning principles into operational controls: role-based training, acceptable use patterns, citation and traceability norms, and lightweight assurance templates for pilots. Include "how we prove it is safe enough" rather than "what we believe".

SUSTAINABILITY AND DECISION THRESHOLDS

Develop shared approaches for measuring and governing environmental impact. This could include procurement questions, usage standards (for example, discouraging low-value prompts at scale), and decision thresholds that link compute intensity to business value and service benefit.

AI-ENABLED OPERATING MODELS AND AGENT GOVERNANCE

Run a themed roundtable on how AI agents change team design, IP boundaries, and accountability. Explore scenarios such as staff onboarding "their" agent tools, and define policy positions on ownership, auditability, and separation of tool, data, and organisational IP.

WORKFORCE REDESIGN AND THE ENTRY-LEVEL GAP

Create a cross-sector program on job redesign, capability pathways, and internships. Focus on preserving learning opportunities, building judgment, and ensuring that automation removes drudgery without removing development pathways.

SOVEREIGN CAPABILITY AND LOCALISATION

Host a discussion on what "sovereign AI" can realistically mean in the Australian public sector context, spanning infrastructure, data, model access, and localisation methods. Emphasise practical options that can be implemented within current constraints.

INNOVATIVE IDEAS AND CASE STUDIES

1. MEASURED PRODUCTIVITY UPLIFT FROM ENTERPRISE COPILOTS

A federal research agency reported full staff access to Microsoft Copilot 365 and early internal reporting indicating an average productivity improvement of 26 minutes per staff member per day. They also implemented an AI transparency statement and responsible use policy before scaling access. This shows a structured pattern: constrain risk early, then expand use with measurement, guidance, and clear accountability.

2. EMBEDDED AGENTS FOR INTERNAL WORKFLOW AUTOMATION

The same organisation described two internally developed agents. One is an offboarding agent in production, automating repeatable steps when staff leave. The second is a skills and capability agent designed to match researchers to projects by drawing together disparate data sources and availability signals. The approach is notable because it targets administrative friction that often slows research delivery, while keeping humans responsible for deployment decisions.

3. FRONTLINE SERVICE INSIGHTS THROUGH SENTIMENT AND SUMMARISATION

A state service organisation described pilots in contact centres using call sentiment analysis and automated summaries to strengthen quality assurance and escalation pathways. This is a pragmatic “low-risk, high-volume” use case: it improves internal responsiveness without automating decisions that affect citizen rights.

4. UNIVERSITY CAPABILITY HUBS BRIDGING RESEARCH AND ENTERPRISE NEEDS

Participants from a major university described standing up an AI hub model that supports cloud-oriented AI services, pre-trained models, and education initiatives (including a cross-university applied AI learning pathway). This kind of “shared capability layer” can reduce duplication across teams, provide safer environments for experimentation, and offer a bridge between cutting-edge research and enterprise deployment.

5. REIMAGINING JUSTICE INTERACTIONS, WITH CAUTION

A participant referenced Nordic-style approaches where citizens can see likely outcomes (for example, indicative penalties) before deciding whether to contest a matter. While presented as a way to reduce unnecessary court interactions and administrative burden, the group also flagged risks around bias, over-reliance, and cultural acceptability. The underlying insight is useful even if the specific model is not directly transferable: AI can enable earlier transparency and self-service, but must be designed with safeguards, context, and clear accountability.

STRATEGIC OUTCOMES AND RECOMMENDATIONS

IMMEDIATE ACTIONS

- **Establish minimum safe-use controls that match real work**: Implement short, role-based training modules that cover sensitive information handling, citation norms, human accountability for outputs, and common failure modes (hallucinations, deepfakes, and overconfidence). Pair this with clear "approved use patterns" so staff know what is encouraged, tolerated, and prohibited.

- **Start measuring value, not just usage**: Move beyond adoption counts. Track time saved, rework reduced, and quality impacts. Use early reference points like "minutes saved per day" to build a baseline, then link those gains to service outcomes or staff experience measures.

- **Create a lightweight assurance pathway for pilots**: Develop a simple checklist for new AI trials: purpose, data handling, risk rating, human oversight, evaluation method, and exit criteria. This helps teams move quickly without losing governance discipline.

MEDIUM-TERM GOALS

- **Translate efficiency into service quality and workforce benefit**: Set expectations that time saved is reinvested into complex work, client support, accessibility, and quality improvement, rather than only output volume. Where possible, tie AI benefits to citizen outcomes (faster resolution times, fewer handoffs, improved accuracy) and workforce indicators (engagement, reduced administrative load).

- **Build shared capability layers to avoid fragmentation**: Where feasible, consolidate tooling into managed environments (for example, internal chat instances or controlled platforms) to reduce shadow use and inconsistent practices. This also supports consistent logging, governance, and learning.

LONG-TERM VISION

- **Treat AI as a core capability, not a sequence of disconnected projects**: Embed AI literacy as a leadership requirement and a workforce expectation, alongside data governance and cyber security. Build organisational "judgement capability" by reinforcing critical thinking, verification habits, and domain expertise.

- **Define a practical sovereign capability position**: Develop a clear stance on what must be sovereign (data, critical models, assurance, localisation) and what can be sourced commercially. This should include an approach to vendor dependency, procurement questions, and options for localised knowledge grounding.

- **Redesign operating models for an agent-enabled workplace**: Anticipate how AI agents change accountability, IP boundaries, and team composition. Establish clear policy positions early so agencies are not forced into ad hoc decisions as agent ecosystems mature.

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