

REASON

GROUP

Safe & Responsible AI in Australia

INDUSTRY SUBMISSION

4 AUGUST 2023

Reason Group outlines a practical initiative to inform the Government on (Target areas Question 7) *How can the Australian Government further support responsible AI practices in its own agencies?*

TABLE OF CONTENTS

INTRODUCTION	3
OBSERVATIONS & RECOMMENDATIONS	4
❖ Leading by Example: Empowering Government with AI for Transparent Decisions, not just regulating	4
❖ Scenario and risk modelling:.....	4
❖ Managing liability:.....	4
❖ Skills development and basic awareness:.....	4
❖ Empower Public Service engagement.....	4
❖ Develop responsible AI as an ecosystem	5
APS-LABS FOR SKILLING AND MODELLING RESPONSIBLE USE OF AI.....	5
Case 1: Real-world whole-of-government problem tackled by government.	5
Case 2: Visualising legislation as a digital twin that facilitates explainability at scale	5
Case 3: Using AI to reverse engineer business rules from data.....	6
THE BUILDING BLOCKS OF AN 'APS-LAB'	6
Three key elements.....	6
Broad applications and benefits.....	7

Introduction

Thank you for the opportunity to respond to the *Safe and Responsible AI in Australia, Discussion Paper, June 2023*.

Specifically, we respond to the *Target area, question 7: How can the Australian Government further support responsible AI practices in its own agencies?*

Reason Group provides some high-level observations, shared experiences, and a practical means to increase the understanding of safe and responsible AI and across the APS within government agencies.

Reason Group is an advisor and capability provider in business and digital transformation to the Australian government. We are also a member of the AIIA, the Tech Council of Australia, and the CSIRO National AI Centre (NAIC) ecosystem that advocates for safeguards and responsible use of AI.

Our response is driven by our firm belief in the importance of a world-class public service that fosters innovation not only within government agencies but also throughout the broader ecosystem and private sector. We have previously submitted insights to the Digital Economy Consultation 2017, the Thodey Review of the APS 2018, and the recent consultation on the initial Data and Digital Government Strategy 2023..

The Australian government's vision to become a leading digital government by 2030, encompassing cyber-safe practices and promoting responsible AI, is both inspiring and commendable. To achieve this, we must address the challenges posed by the increasing accessibility of Generative AI and technologies like ChatGPT since its debut in November 2022.

Responsible AI practices entail not only ensuring safety but also emphasizing ethical use, diversity, inclusion, morality, cultural harmony, and governance of the technology and the data it utilizes. While AI presents boundless possibilities, it also raises complex dilemmas daily. The initial Data and Digital Government Strategy (2023) acknowledged the role of AI/ML but only in passing, underscoring the urgency for a comprehensive approach that places AI at the forefront.

No one can predict what human intelligence is capable of achieving. However, we do know that artificial intelligence has the power to accelerate, amplify and scale it. This means that as humans and systems are imperfect, will the imperfections accelerate or be attenuated by AI?

Regulation has struggled to keep pace with AI development, as evident from the 2015 open letter on artificial intelligence calling for a pause to allow regulatory catch-up. Finding system-based solutions is crucial to putting AI regulation on the front foot. This requires not only catching up with regulation but also fostering readiness, extensive knowledge, and in-depth research into the societal impacts and benefits of AI.

A key principle of AI regulation should be to ensure that humans remain in control. It is imperative to prevent AI from being controlled by uncontrollable actors or unstable systems

As we integrate AI into our "GovTech," it becomes even more vital to develop a forward-looking strategy that positions Australia as a leading digital economy and society by 2030.

We hope these observations resonate and provoke action in safe and responsible AI, tailored to the needs of people and businesses, the APS, staff, and multi-jurisdictional government partners .

Observations & recommendations

We support regulation of the safe and responsible use of AI, but the speed of policy and regulatory agility needs a step change and new approach – the solution to which may draw on AI technology itself. In the context of ‘existential threats’, there is much discussion on safeguards to protect people from harm, but that term connotes abuse and intentional misuse when a larger threat looms in the form of inappropriate use or incompetent use that the ‘humans in the loop’ may not be aware of.

- ❖ **Leading by Example: Empowering Government with AI for Transparent Decisions, not just regulating:** Government agencies are faced with an unprecedented number of decisions to make daily. With the advent of AI, these capabilities are scaling up rapidly. However, when it comes to decisions that impact citizens or businesses, the process can seem like a black box - complex data factors, business rules, and legislation all play a role, leaving individuals with little explanation on the final outcome. AI can provide much-needed clarity at scale, ensuring that decisions are not only accurate but also transparent. Imagine a future where AI enables decision explainability, revolutionising the way government operates. From granting or removing benefits to handling various administrative tasks, AI will be an integral part of the government architecture. Decades worth of data, intricate legislation, regulations, codes, and ever-changing rules for eligibility and benefits will all be efficiently administered with the aid of AI. The key to success lies in unifying data, digital technology, AI/ADM (Autonomous Decision-Making), and cybersecurity under one cohesive strategy and governance control. By establishing a common whole-of-government governance framework, we can ensure that AI is seamlessly integrated into the heart of government operations. The ultimate goal is to empower the government to lead by example, not merely through regulation but through concrete action. By leveraging AI responsibly and ethically, the government can set a new standard for transparent decision-making, setting the tone for other sectors to follow suit. Let's envision a future where AI is the driving force behind efficient, accountable, and people-centric government processes.
- ❖ **Scenario and risk modelling:** Core legislation should be reviewed to identify all the decisions the legislation entails and gives authority to. Decisions can then be analysed and simulated against a risk-based AI framework that makes it clear what is relatively high risk and what is low risk; what are the decision-making use cases; where is the use of AI responsible, and what areas should be quarantined from AI.
- ❖ **Managing liability:** Government may also need to play a role in explaining the position of AI accountability as well as simplify the insurance relief for (small and medium) businesses that already carry insurance for data protection, and cyber and many other policies that is already growing in cost and complexity.
- ❖ **Skills development and basic awareness:** Government invests heavily in human and system capability uplift, this training in many supporting disciplines such as security and privacy, risk, quality and so on. All those elements are driven by recognised protocols and standards that we must draw on when we design, build, deploy and use AI-enhanced systems and processes at the same time as understanding how the standards in AI are evolving in order to stay current with those guardrails and make those standards lived-in practices.
- ❖ **Empower Public Service engagement:** Anything will look like it is moving too fast when we are standing still. Government agencies do not have the right platforms and environments for the APS to understand and experience AI collaboratively and broadly through a hands-on environment that makes it safe to work with data, foundation models, and algorithms to explore applications for innovation, model the risks, and simulate test cases. APS needs ‘sandpits’ to share best practices, e.g. prompt libraries and standards for reuse, and adaption to speed up responsible use. These safe house environments are what is needed for

more cross-agency and cross-jurisdiction experimentation. There is much value to be gained by sharing best practices from industry, as well as defence and security, academic researchers and technology start-ups, with public sector administration more broadly.

- ❖ **Develop responsible AI as an ecosystem:** The APS Modernisation Fund and the proposed Digital Resilience Fund are sources of funds for responsible use of AI sandpits aligned with the vision of ecosystem approaches to deliver government services. We also support mechanisms like industry grants and the National Reconstruction Fund in the government's investment strategy.

APS-Labs for skilling and modelling responsible use of AI

This section offers an example from recent work we have undertaken with the Australian Government in three areas that we think illustrate how we learn more about the responsible use of AI in an integrated way.

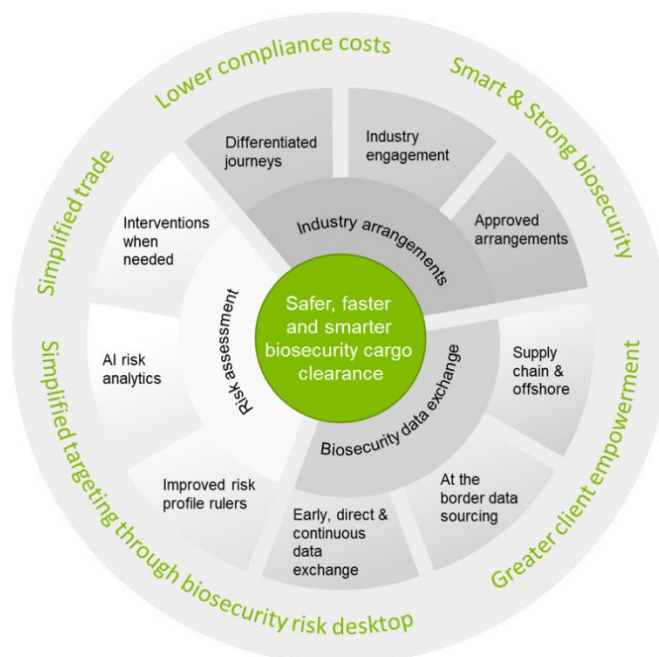
Case 1: Real-world whole-of-government problem tackled by government.

The Simplified Targeting and Enhanced Processing System (STEPS) Program was initiated by government to improve the import continuum through safer, faster and smarter biosecurity cargo clearance. STEPS is a key capability for the Simplified Trading System (STS), Digital Economy strategy and the Deregulation agenda.

Our innovative approach for this government-with-industry engagement was based on safe data sharing, simulated technology environments, and accelerated co-development within 13 weeks to deliver a breakthrough to trade data exchange solution between government and industry to alleviate supply chain bottlenecks that the economy suffered throughout the global pandemic.

We used AI to help visualise the trade flow data that existed in multiple systems across government and industry. Visualising the data helped clarify the possibilities, practicalities and the risks of sharing data in a safe and responsible way across:

- ❖ Continuous Biosecurity Risk Assessment to test the feasibility of integrating trade and biosecurity data continually as goods move through the system, moving to a model where we assess the risk is assessed using AI at multiple points along the import supply chain.
- ❖ Supply Chain Data Exchange to develop a data exchange capability to provide a secure platform to integrate with external data providers across the supply chain continuum.
- ❖ Enhanced Business Partnerships that explored the business model drivers that enabled effective ecosystem integration and collaboration.



Case 2: Visualising legislation as a digital twin that facilitates explainability at scale

With an AI technology partner, we ran a trial that used AI Natural Language Processing (NLP) to analyse a large piece of legislation and related instruments that are constantly updated for a very large agency to allow

administrators and experts to regain control of their business rules – dynamically in a digital context to enable true policy agility by decoupling business rules from IT systems.

The visualisation is a digital twin of the legislation, a tool for business policy experts who understand legislation to get total control over their rules and move forward on AI/ADM decision explainability against the legislation.

The model could then be used to feed decision-making open via API for chatbots, M2M, and workflow engines where complex rules and policy changes could be tested in a safe quarantined environment that is used as a legislation analytics service lab, as well as make explainability as personalised customer-centric form rather than legislation-centric form.

Case 3: Using AI to reverse engineer business rules from data

We are often involved in government legacy maintenance, modernisation, and replacement. One of the largest problems is the detailed current state of the rules and how the patchwork of systems and manual processes come together. That current state is not documented and would not be able to be described with accuracy even through armies of business and systems analysts.

An innovative approach taken by another of our AI-partner is to take a 'data-first' approach. This means applying AI to analyse the source data, inputs, integrations, data transformations and outputs in order to derive business rules which could then augment the work of business ad system analysts. To perform this exercise, a safe quarantined environment is also necessary to model the responsible use of AI to discover hidden business rules across legacy systems.

The building blocks of an 'APS-Lab'

Both federal and state governments are encouraging more use of proofs-of-concept in the procurement of ICT capability and solution innovation. We also believe this is the smart way to uplift APS knowledge, capability and test the responsible use of AI within government agencies.

Three key elements

The elements of the approach, the patterns, development cadence and the safe technical environment we used can be applied broadly across government as follows:

- ❖ **Technical sandpit** - a pre-configured turn-key Microsoft Azure as-a-Service virtual environment that emulates operating environments that allow the agencies to work effectively with industry partners. The technical components include government foundation technology that reflects the Australian Whole-of-Government Architecture and capability model that includes emerging technology and AI-enabled tools. The sandpit is a safe sovereign environment that is PSPF and ISM compliant. This capability can be provided as a standard managed service offering that eliminates the time and hassle of agencies trying to stand-up and manage temporary facilities, multi-vendors and licences.
- ❖ **Hyper agile development programs and methodology** and a team of small, medium, or large multi-disciplinary team configurations, working to a cadence to iterate working software and data software over 13 weeks – not months. The 13-week program orchestrates co-development cycles, ingests data mirrored from production systems that can be shared in a safe environment, and real-time AI-enabled data visualisation of policy, rules, loads, and risks from use cases. The process compacts discovery and designs

which speeds insights and proves feasibility within the fixed time and cost that would otherwise produce simplistic outputs from post-it note sessions.

- ❖ **Right people** from the ecosystem are essential to the MDTs. Participants from government included SMEs. The team must reflect a microcosm of the diversity of actors working with experts in policy, architecture, data science, solution design, platforms and complex integrations. Rapid expert feedback loops to sustain high-value iterative development.

Broad applications and benefits

The APS-Labs approach is a viable turn-key approach for cross-agency, cross-jurisdiction and ecosystem co-development that can be easily replicated and beneficial at any stage of a data and digital initiative to help de-risk the implications of major commitment decisions. They are safe houses to incubate ideas in simulated environments, that can then be transferred in-house, and then used to substantiate potential for reuse.

Government ICT strategies and digital initiative roadmaps are looking into the future but are often unaware of what each other is doing or planning. APS-Labs are places to ideate government architecture of the future and make the longer-term architectures visible as landing zones that government CIOs, CDOs, and CISOs can factor into their data and digital plans. And in so doing improve the collective quality of our short-term and 10-year planning horizons.

Another valuable application is to assure delivery in business cases for large and complex ICT by testing the options and design approaches, costing models, extracting real risks and issues, and prioritising high-value business requirements.

We see several scenarios for the APS-Labs approach using testbeds and sandbox environments:

- **Green-fields** – as a testing ground for proofs-of-concept for new models in an unintegrated mode (similar to the approach used by STEPS above)
- **Blue-fields** – hybrid testing of existing tech with new tech, testing solution viability to progress into business cases and full integration
- **Brown-fields** – replicating the whole of government architecture for proving reuse and integration of new solutions/products to show a “path to production” (particularly for legacy modernisation)
- **Sandboxes** – testing of live alpha/beta products and services (similar to the sandbox approach used by ASIC for FinTech and RegTech)

As AI/ADM capabilities unfold, the APS-Labs are needed more than ever as safe places to test use cases with broad involvement to achieve rapid minimal viable understanding and to promulgate responsible AI practice with targeted knowledge transfer.

About Reason Group

Reason Group is a homegrown Australian business and technology solutions firm that specialises in digital government.

Our purpose is to make government easier for those that deliver it and the people it serves. We partner with you to unlock opportunities within the unique complexity of our Australian government landscape.

Our team brings both hearts and smarts to the table. We are diverse and deeply experienced and to put it plainly, we care. We don't play around the edges; we get to the heart of the problem to ensure we're in it for the right reasons. We want to know what makes you tick and how we can make an impact.

Our approach is personal. We tailor our methodology to every client's individual needs, rather than fitting your needs to our approach. We embed ourselves seamlessly in your space, to see projects how you see them. Success is achieved through empathy, collaboration, and a commitment to delivering no matter what.

Reason Group Pty Ltd

Level 11, 68 Northbourne Avenue

Canberra ACT 2601

ABN 34 128 711 348

T +61 (0)2 6152 0942

business@reason.com.au

www.reason.com.au

This document is prepared solely for the intended recipient and client personnel. The material presented reflects Reason Group's best judgement based on the available information at the time of preparation. This document contains confidential and proprietary information and remains the property of Reason Group. No part of this document may be used, circulated, quoted, or reproduced for distribution outside the client organisation without prior written approval from Reason Group.

© 2023 Reason Group



making government easier

