AI in the Public Service

AI is like a powerful engine. But an engine needs more than just fuel to run well: It needs a steering system and rules to guide its path on the road. In New Zealand's public service, steering and guidance comes from a mix of frameworks and governance efforts. Together, they ensure that AI is used safely, responsibly and reflects our public values and earns people's trust. The Public Services Commissioner appointed the New Zealand Government Chief Digital Officer (GCDO) as Public Service AI Lead.

GCDO has been asked to support safe and responsible Public Service AI use including coordinating system leadership and standard setting, stewarding system AI adoption, supporting agencies to use AI and working with providers to drive value for money in the services they provide us.

Public Service AI Framework

This starts with the Public Service AI Framework: A foundational guide that lays out a shared vision for how AI should be used across the public service. It helps ensure that our AI is ethical, transparent, inclusive and aligned with our public service values.

There are five key principles, adapted from global best practice like the OECD AI Principles. Cabinet agreed in June 2024 to promote the OECD principles as a key direction for our approach to responsible AI in New Zealand. The OECD principles promote the use of AI that's innovative and trustworthy and that respects human rights and democratic values.

Here is more detail and a description for each of the five principles:

Inclusive, sustainable development

- Innovation

- Efficiency

- Resilience

- Address inequalities and transition equitably

In designing a public service AI tool, the team considers which communities would be impacted most directly by our AI solution and engaged them in consultation during the design phase for our model.

By considering which communities are most impacted by our AI tool and consulting them, we make sure we develop our tools inclusively and in line with community needs.

Human-centred values

- Privacy and data ethics

- Upholding democracy and the rule of law

- Human and labour rights

- Human oversight

A healthcare AI tool prioritises patient autonomy by offering personalised treatment options while giving the user the option to make changes that they see fit to their personal circumstances.

AI and other tech tools are effective ways to offer bespoke treatment, but the patient is always at the centre of the decision-making process and what care is best for them, including the way they use AI tools to improve their care.

Transparency and explainability

- Promote awareness and understanding of AI systems and interactions.

- Explainable outcomes

A financial AI tool provides users with a detailed breakdown of how investment recommendations are generated, including data sources, written in a way that makes intuitive sense to the user.

We are showing where we sourced the data so that the user can check for themselves, in plain language that a layperson can understand.

Safety and security

- Safety by design

- Data protection

- Traceable data

- Robust risk management, including national security risks

A Generative AI chatbot is made directly available to the public to answer questions about social benefits – any personal information disclosed by the user is not stored.

The chatbot prioritises the privacy of the user and protects its data, so that the user can understand the mechanism behind the decision and in turn, fosters greater trust in the system they're using.

Accountability

- Governance

- Regulatory frameworks.

- Auditing with human oversight.

- Capability growth

A public service AI tool has a dedicated team responsible for monitoring system performance and addressing user concerns promptly to ensure accountability.

The dedicated team can track and address any concerns with our AI tool, ensuring we can audit and amend the tool with human oversight.

GCDO Public Service AI Work Programme

To support agencies applying these principles, the Government Chief Digital Officer (GCDO) leads a public service AI work programme. This work programme is structured around six pillars:

1. Governance — to support transparency and human accountability in Public Service AI use.
2. Guardrails — to enable safe and responsible Public Service AI use.
3. Capability — to build internal and external AI knowledge and skills.
4. Innovation — to provide pathways to enable safe AI testing and innovation in the Public Service.
5. Social licence — to ensure New Zealanders and public service workers have trust and confidence in Public Service AI use.
6. Global voice — to ensure international counterparts see New Zealand as a trusted AI partner.

Below are some everyday examples of public servants encountering the pillars of the AI Work Programme in their day-to-day work.

1. Governance: "One of our team members has a great idea for a new AI tool and has put forward a use case to the steering group at the end of the month. They will decide whether or not it meets our privacy and security guidelines to be approved".
2. Guardrails: "I want to use a Gen AI tool to analyse some data for me, but I am not sure whether I can upload it to an online generative AI tool . I'll have a chat with my manager to ask whether or not there's an approved tool we can use. Additionally, I will ensure that the AI policy is reviewed and any guidance from the GCDO is also considered, beyond consulting with my manager".
3. Capability: "One of our team members is going to host a monthly Lunch and Learn session on AI. Our next session is about an exciting new tool called ABC Search and how it can help with in-depth research".
4. Innovation: "Our department has set up a sandbox environment where we can safely experiment with new AI models. This allows us to test innovative solutions without impacting live data or services, ensuring any potential risks are managed before full deployment."
5. Social license: "We've organised a community forum to gather feedback from the public on our use of AI in government services. This is to ensure transparency and build trust by addressing any concerns they may have about privacy and ethical use."
6. Global voice: "Our team is preparing a presentation for an international AI conference to showcase New Zealand's advancements in ethical AI governance. By sharing our practices and successes, we aim to strengthen our reputation as a leader in safe and responsible AI use."

Key legislation and its influence on AI use

New Zealand's approach to our regulation is to rely first on the regulations and frameworks that are already in place. We do not have AI-specific legislation in place like our counterparts in the EU and elsewhere. This means that as New Zealand public servants interested in AI, there is more than one home for the "rules" that govern how and when we use AI . The upcoming AI Assurance Model will guide agencies through a single framework.

This section will provide a high-level summary of the rules of the road. You do not need to know each of these in excruciating detail, but it is helpful for you to understand the direction of the public service as it concerns AI.

The Treaty of Waitangi

As New Zealand's founding document, te Tiriti key to a lot of the work we do as public servants, working together to strengthen the Māori-Crown relationship. There are a number of implications on data, artificial intelligence and similar kaupapa that will be covered in further detail in Module Two.

Official Information Act 1982

This Act facilitates transparency by ensuring the public can access government information. When it comes to AI and data, we want to be accountable to the public about the AI systems we use, to ensure the public are informed about how we use AI.

New Zealand Human Rights Act (1993)

This Act prohibits discrimination and promotes equal treatment. In the context of AI, it means that AI solutions must be designed and implemented in ways that actively avoid bias and ensure fair treatment of all of those who use the tool.

Bill of Rights Act 1990

This Act safeguards fundamental human rights and freedoms. AI systems must respect those rights, such as privacy and freedom from discrimination, guiding how we ethically deploy AI in delivering public services.

Copyright Act 1994

Creators have rights over the works they create. This Act has implications on intellectual property rights and fostering innovation within legal boundaries.

Public Records Act 2005

The Act is all about how government records are created, maintained and preserved. Proper data management practices are key, making sure that our AI decisions can be audited and that we maintain data integrity over time.

Public Service Act 2020

The legislation that governs a unified and responsive public service, we need to integrate AI technologies in ways that enhance service delivery for the end user and in line with upholding our public service values.

Privacy Act 2020

It is critical that you consider how privacy is protected when personal data is involved. This Act governs how we collect, use and disclose personal information. AI tools need to mandate responsible data handling practices, ensuring that we deploy AI in a way that builds public trust through transparency of how we use data.

Algorithm Charter for Aotearoa New Zealand (2020)

The charter states that humans must review and decide on almost all significant decisions made by government agencies. Significant decisions are those where people's wellbeing is impacted, or there is a high likelihood many people could suffer an unintended adverse impact, such as incorrect tax decisions. Additionally, the charter requires transparency of algorithm use by keeping appropriate records and providing clear and accessible documentation.

This is not an exhaustive list, but it is an introduction to the key pieces of legislation and other regulations that govern how, when and why we can use AI in the New Zealand public service.

Let's consider an example.

Imagine your agency is implementing an AI tool to help analyse large volumes of public feedback. The tool automatically scans and categorises open-text submissions from individuals and communities, identifying key themes and flagging potentially harmful content. It stores submissions securely, does not make any final decisions, and all outputs are reviewed by human staff before use.

Below you can see how the legislation and regulations might apply.

Privacy Act - Since the tool handles comments that may include personal information, privacy protections are likely required.

Official Information Act - If someone requests information about how their feedback was used, the agency may need to provide an explanation.

Public Records Act - If the outputs or summaries are used in decision-making, they may need to be stored appropriately.

Human Rights Act - If the tool affects some groups differently, this legislation could become relevant.

Bill of Rights Act - If individuals' views are treated differently or if content is removed, this Act may apply.

Te Tiriti o Waitangi - If Māori perspectives are included, it is important to consider cultural relevance and equity.

Public Service Act - This supports integrity and accountability in public service, which are essential when using emerging technologies.

Algorithm Charter - This tool is likely covered by the Charter, which encourages agencies to explain how algorithms work and ensure fairness.

Copyright Act - This may be relevant if comments are reused in published materials or training content.

There may be specific things you are working on in your team and at your agency that will mean you must have regard to other legislative or regulatory frameworks. However, it is valuable for you to be familiar, at a high level, with New Zealand's approach to AI regulation, which Cabinet agreed is focused broadly on:

1.   A light-touch, proportionate and risk-based approach.

2. Recognises our policy context already providing some guardrails for AI use.

3. Further regulatory intervention needs to unlock innovation while addressing acute risks.