Good Ancestors
NEW ZEALAND

Recommendation: Uplift government AI safety and security capability

Key points

- The Five Eyes cyber security agencies warned in June 2026 that leading AI systems are rapidly transforming cyber risk.
- New Zealand's security and policy agencies do not yet have the capability to match the pace of that change, and no agency has a clear mandate to assess frontier AI capability and advise Ministers on it.
- We recommend investing to increase the AI security capabilities of the National Cyber Security Centre and the Electoral Commission, together with AI risk capability across government.

Why it matters

AI is being used to write malicious code, find security weaknesses, create deepfakes, clone voices and scam New Zealanders. That puts critical infrastructure, government services, election integrity, personal data and businesses at risk. A 2021 ransomware attack took Waikato hospital systems offline for weeks and attacks in 2020 disrupted NZX trading for four days. Leading AI systems are making such attacks cheaper and easier.

In June 2026 the Five Eyes cyber security agencies, including New Zealand's National Cyber Security Centre (NCSC), issued a joint call to action warning that leading AI models can identify and exploit vulnerabilities at unprecedented speed and scale. The NCSC responded to 1,164 incidents in the first quarter of 2026, and estimated the cost of online threats to New Zealanders in 2024 at $1.6 billion.

What we're proposing

We recommend targeted increases in funding and capability across the New Zealand government:

- **National Cyber Security Centre**: Additional funding for AI-specific threat detection and response, and for the technical advice it gives agencies and critical infrastructure operators. The NCSC has begun publishing guidance on securing agentic AI systems, and that work needs resourcing to keep pace with the rapidly advancing AI frontier.
- **Electoral Commission**: Funding for detecting and countering AI-generated disinformation, and to strengthen its wider defences against emerging AI threats to electoral integrity over the next three years. The Commission has said it is confident in the security of its systems, but given the rapid pace of AI advancement, those defences will need continual review, and work should start now if we are to address the risk of widespread disinformation affecting the 2029 election.
- **AI risk policy capability:** Dedicated policy staff to advise the government on AI safety, based in the Ministry of Business, Innovation and Employment and the Department of the Prime Minister and Cabinet. The government should also name the agency accountable for assessing frontier AI capability and advising Ministers, or establish an AI Safety Institute to do it.

This works alongside our proposals for an AI Safety Strategy and an AI Safety Institute: the strategy sets the plan, the institute provides expert assessment, and this funding builds the capability to act on both and advise ministers across the broad societal implications of AI development.

The opportunity for New Zealand

Building AI capability in multiple agencies would equip the government to respond to a rapidly evolving threat environment that spans private business, government services and individual New Zealanders.

This capability could also strengthen New Zealand's role in the Five Eyes security alliance, equipping it with the technical depth to engage deeply with allies and act on the warnings it receives.

Strong AI capability in government supports public confidence in AI systems, and that confidence is a precondition for widespread adoption. Without it, companies may hesitate to put AI into payroll, patient records or payment systems, even where it would save customers money, for fear of backlash. Creating the conditions for adoption starts with the government understanding AI and building policy around it.

How it could be done

The proposal can be delivered through existing agencies and normal Budget processes. It is a funding and staffing decision: no new institution, legislation or obligations on business would be needed.

The NCSC and Electoral Commission increases could be sought through Budget 2027 and appropriated within existing votes. Expanding AI policy capability could begin sooner, since departments can hire or reassign staff within current operations, while longer-term funding is settled.

International precedent

New Zealand was the last OECD country to publish an AI strategy. The United Kingdom's AI Security Institute has around £100 million in committed funding and more than 30 technical staff working alongside the UK's National Cyber Security Centre.

Australia has centralised its AI capability in an Office of AI within the Department of the Prime Minister and Cabinet, established an AI Safety Institute, and its Electoral Commission runs a disinformation register.

Australia also spends far more per capita than New Zealand on its intelligence capability. The Australian Signals Directorate, which covers communications security, intelligence and cyber security, has an annual budget of about AU$2.5 billion (close to $3 billion). In contrast, the Government Communications Security Bureau (GCSB) has an annual budget of only $344 million.[1] Even allowing for Australia's larger population, that is roughly 70 percent more per person.

This proposal would align with the approach to AI safety New Zealand has endorsed in international forums. In 2024, it signed the Bletchley Declaration and the Seoul Ministerial Statement, both of which included an international agenda for managing frontier AI risk that included building public sector capability.

Frequently asked questions

Didn't the government just increase this funding? Budget 2026 provided a $156 million uplift for the intelligence and security agencies. It is not yet clear how that will be allocated, and New Zealand would still spend far less per person than Australia. Increases are needed specifically in AI capability, election security and AI policy capacity.

Isn't this a job for the private sector? Firms can secure their own systems, but they cannot assess frontier model capability or defend an election. Both require government capability, and small businesses depend on the NCSC for advice they cannot produce themselves.

How much would it cost? In the order of $15–25 million per year. This could increase the capacity of the NCSC, Electoral Commission and MBIE to employ policy staff across multiple areas of AI security and safety. That is roughly 1 percent of what online threats cost New Zealanders each year.

[1] GCSB's annual budget is $344 million in 2026/27 and the Australian Signals Directorate's annual budget is approximately $2.5 billion AUD.