19 August 2025

s 9(2)(a)

Official information request
Our ref: R001037

Tēnā kōrua

Thank you for your Official Information Act 1982 (OIA) request received by the Ministry for Regulation (Ministry) on 6 June 2025. On 3 July 2025, we notified you of the extension to make our decision to 15 August 2025.

We have itemised your request for ease of reference. You requested:

> *"Please provide policy documents and discussions involving Ministry of Regulation officials relating to the use of AI in processing of submissions on the Regulatory Standards Bill. If third party providers such as Public Voice, or other third party providers were used please provide communication between the Ministry and those third party providers as related to the use of AI and in particular, any discussions or policy documents relating to the following issues:*
>
> 1. *Was the data processed in any country other than NZ?*
> 2. *What privacy guarantees were made regarding the use of models to process this data?*
> 3. *Was the ministry assured that the AI model provider would not be able to train and improve their models using the data uploaded, and if so how much confidence did the ministry have in these assurances?*
> 4. *Was the provider or providers required to delete any data uploaded after processing?*
> 5. *Was the data redacted for privacy before it was uploaded to any AI models?*
> 6. *If so, who did the redaction work, what instructions were they given and what types of data was redacted?*
> 7. *If not, what representations did the third party providers or AI model providers make in regards to data redactions, if any?*
> 8. *What consideration, if any, was made regarding use of indigenous data and data provided Māori?*
> 9. *Was any of the data provided to the model written in te reo Māori and if so, what consideration was taken for safeguarding the use of that data and preventing it being used to train language models?*
> 10. *Finally, could you please provide any privacy impact assessments that were done in relation to this work."*

Response
I have responded to each of your questions in turn.

1. Was the data processed in any country other than NZ?
Data processing occurred both within and outside of New Zealand. Please also refer to the response to **item 2**.

2. What privacy guarantees were made regarding the use of models to process this data?
The Ministry worked with PublicVoice to ensure the privacy and security of all submissions. Privacy implications and risks were considered appropriately and the Ministry sought and received assurances regarding the handing of personal information, including:

- Confirmation that all information would be kept confidential, in line with PublicVoice's privacy policy and in line with legal requirements under the Privacy Act 2020.
- Confirmation that data would not be used to train Artificial Intelligence (AI) models and assurance that all data would be destroyed upon project completion.
- Zero data retention policies were in place, and SSL/TLS 1.3 encryption was used for data transmission.
- Data would be processed in New Zealand, the European Union or the United States, and only three New Zealand-based staff would have access to the information.
- Outputs would be human-reviewed, and a multi-step validation process would be used to mitigate risks such as hallucination.

3. Was the ministry assured that the AI model provider would not be able to train and improve their models using the data uploaded, and if so how much confidence did the ministry have in these assurances?
Yes, the Ministry received and accepted the assurances on this matter. Please refer to the response to **item 2**.

4. Was the provider or providers required to delete any data uploaded after processing?
Yes, the tools used to analyse the submissions did not retain personal information or submission content as the information was destroyed upon project completion. Please refer to the response to **item 2**.

5. Was the data redacted for privacy before it was uploaded to any AI models?
Personal information included in the submissions were analysed by AI technology as part of the analysis of all submissions. The tools used did not retain personal information or submission content as the information was destroyed upon project completion.

Submitter identification data was analysed as part of data handling and cleaning processes for the following purposes:

- Classification of each submission by submitter type (individual, iwi/hapū, or organisation)
- Determination of stance (oppose, support, partially support, or unclear)
- De-duplication and flagging of identical or near-identical submissions
- Determination whether submissions contained an Official Information Act request
- Identifying the language and whether a submission needed translation

- Assessment whether submissions pertained to matters unrelated to the proposed Regulatory Standards Bill.

6. If so, who did the redaction work, what instructions were they given and what types of data was redacted?

Refer to the response to **item 5**.

7. If not, what representations did the third party providers or AI model providers make in regards to data redactions, if any?

There were no commitments specifically with regard to data redactions. However, contractual obligations on PublicVoice included the following privacy requirements:

- To comply with the Ministry's Privacy policies and guidelines
- To comply with the Privacy Act (and other laws)
- To access the information only to the extent necessary to provide the services
- To not disclose the information to any third party
- To keep the information secure
- To return or destroy confidential information after expiry or termination of the agreement.

Beyond the above, the Ministry specifically tested and received reassurances on further potential privacy impacts, including:

- No information would be used to train the AI model
- Information would be destroyed after expiry or termination
- The number of people with access to the information would be limited (3), and in New Zealand
- Measures to take to mitigate risk of AI error/hallucination
- The information will be kept confidential, for both privacy and general confidentiality purposes.

8. What consideration, if any, was made regarding use of indigenous data and data provided Māori?

We understand your question to be seeking clarification on how the Ministry managed data and information from submissions provided by Māori. These submissions were handled with care and integrity, consistent with the approach taken for all submissions. The Ministry maintained a record of submissions received in te reo Māori, which were subsequently translated to ensure accessibility and understanding. To support this, PublicVoice was asked to run a script to identify any additional submissions requiring translation. The translation work was undertaken by the Department of Internal Affairs.

9. Was any of the data provided to the model written in te reo Māori and if so, what consideration was taken for safeguarding the use of that data and preventing it being used to train language models?

Refer to the response to **items 2, 5, 7** and **8**.

The Ministry was assured that any data provided was subject to the data handling and cleaning processes referenced in **item 5**, and we were advised by PublicVoice that data would not be

retained/used for training AI models as per **item 7**. Please also refer to the privacy assurances in **item 2** and the treatment of submissions in te reo Māori in **item 8**.

10. Please provide any privacy impact assessments that were done in relation to this work.
While the Ministry does not hold a 'Privacy Impact Assessment' document in the form of a standalone report, privacy implications were considered appropriately and the Ministry sought and received assurances from PublicVoice regarding the handling of personal information as part of its due diligence as explained in the response to **item 2**.

Documents
You requested policy documents, discussions and communications involving Ministry officials and PublicVoice relating to the use of AI in processing submissions on the discussion document.

On 29 January 2025, Ministry officials discussed the potential use of AI during a weekly meeting with the Minister. At the meeting, we advised we were exploring engaging with PublicVoice and the application of AI tools for the analysis of submissions. The Minister was supportive of this.

We provided briefing papers to the Minister for Regulation, which reference working with PublicVoice and the use of AI in the analysis of submissions. Copies of these briefings are publicly available on the *Publications and resources* page of the Ministry's website, I therefore refuse the publicly available parts of your request under section 18(d) of the OIA as the documents you have requested, itemised below, are publicly available.

- MFR2025-026: Regulatory Standards Bill Initial findings from public consultation[1], dated 21 February 2025
- MFR2025-027: Regulatory Standards Bill Summary of Submissions[2], dated 19 March 2025
- Information Release - Summary of Submissions for proposed Regulatory Standards Bill[3], dated May 2025.

Some information in the aforementioned material has been redacted consistent with the provisions for withholding information under the OIA. Where this is the case, the relevant sections of the OIA that would apply have been identified and where information was withheld, no public interest considerations were identified that would outweigh the reasons for withholding the information.

I have considered the grounds under which information has been redacted in the proactively released documents which you have requested, and I consider they continue to apply in the same ways under this request. I therefore withhold the same parts of these documents, under the same grounds as listed in the published versions.

[1] https://www.regulation.govt.nz/about-us/our-publications/mfr2025-026-regulatory-standards-bill-initial-findings-from-public-consultation/
[2] https://www.regulation.govt.nz/about-us/our-publications/mfr2025-027-regulatory-standards-bill-final-summary-of-submissions/
[3] https://www.regulation.govt.nz/about-us/our-publications/information-release-summary-of-submissions-for-proposed-regulatory-standards-bill/

Communications involving Ministry officials and PublicVoice, which are in scope of your request are attached as **Appendix A**. Some information has been withheld under the following sections of the OIA:

- 9(2)(a) to protect the privacy of natural persons
- 9(2)(g)(i) to maintain the effective conduct of public affairs through the free and frank expression of opinions by or between or to Ministers of the Crown or members of an organisation or officers and employees of any public service agency or organisation in the course of their duty
- 9(2)(h) to maintain legal professional privilege.

As required by section 9(1) of the OIA, I have considered whether the grounds for withholding the information requested is outweighed by the public interest. In this instance I do not consider that to be the case.

Additional information

Attached as **Appendix B** is a copy of the Ministry's Artificial Intelligence (AI) policy, which sets out how AI may be used within the Ministry. This policy explicitly stipulates that Ministry staff may use AI tools to analyse and summarise Ministry information, including submissions data. While the material is not strictly within the scope of your request — as the policy was confirmed and came into effect on 29 April 2025 — we have decided to provide this to you for your reference.

Right of review

If you wish to discuss this decision with us, please contact hello@regulation.govt.nz.

You have the right to seek an investigation and review by the Ombudsman of this decision. Information about how to make a complaint is available at www.ombudsman.parliament.nz or freephone 0800 802 602.

Please note that we may publish this response (with your details removed) on the Ministry for Regulation website.

Ngā mihi

s 9(2)(a)

Aisling Risdon
Head of Ministerial Services
Ministry for Regulation

Appendix A

From: Adam Jackson
Sent: Wednesday, 29 January 2025 11:18 am
To: Pip Van Der Scheer
Cc: Isabelle Sin
Subject: Privacy

Hey, circling back on privacy. Sounds like it is very low risk. It's covered in the contract we're using and s 9(2)(h) of course, we were expecting them to comply with their obligations under the Privacy Act and contract and keep personal information confidential and appropriately secure.

Let me know when you think we have our ducks in a row enough for a further conversation with the provider and I'll line that up.

Adam Jackson
Chief Advisor to the Chief Executive
Ministry for Regulation
Mobile: s 9(2)(a) **| Email:** adam.jackson@regulation.govt.nz

Ministry for Regulation
Te Manatū Waeture

www.regulation.govt.nz

From: s 9(2)(a) @publicvoice.co.nz>
Sent: Thursday, 30 January 2025 1:39 pm
To: Adam Jackson; Pip Van Der Scheer; Isabelle Sin
Subject: RE: RSB Submission Analysis

Hi All,

Good to catch up. Some more information about some of AI tools we use to improve productivity.

MAXQDA (PublicVoice is a certified reseller in New Zealand) is one of the AI tools we use to speed the analysis process up. Here is some more information about the security - AI data protection - MAXQDA and AI Assist for qualitative data analysis - MAXQDA

Most importantly

- No AI model training on user data
- SSL/TLS 1.3 encryption for data transmission
- Zero data retention

Cheers,
s 9(2)(a)

-----Original Appointment-----
From: Adam Jackson <Adam.Jackson@regulation.govt.nz>
Sent: Thursday, 30 January 2025 10:02 am
To: Adam Jackson; s 9(2)(a) ; Pip Van Der Scheer; Isabelle Sin
Subject: RSB Submission Analysis
When: Thursday, 30 January 2025 12:30 pm-1:00 pm (UTC+12:00) Auckland, Wellington.
Where: Microsoft Teams Meeting

Hi s 9(2)(a) ,

I've had something come up between 12 and 12.30 so let's grab the second half of that slot. Talk soon.

Thanks
Adam

Microsoft Teams Need help?

Join the meeting now

Meeting ID: 410 041 277 41

Passcode: K5ni6YX2

Dial in by phone

+64 4 889 8046,,399072101# New Zealand, Wellington

Find a local number

Phone conference ID: 399 072 101#

For organizers: Meeting options | Reset dial-in PIN

From: Aimee Riddell
Sent: Thursday, 30 January 2025 3:26 pm
To: Jeremy Shoebridge; Adam Jackson
Cc: Pip Van Der Scheer; Laura Fair; Isabelle Sin
Subject: RE: AI usage by provider

Thanks Adam, nothing further to add at this point.

Agree s 9(2)(h) retention and disposal after completing work, suggest adding under item (a)

Cheers,
Aimee

From: Jeremy Shoebridge <Jeremy.Shoebridge@regulation.govt.nz>
Sent: Thursday, 30 January 2025 3:19 pm
To: Adam Jackson <Adam.Jackson@regulation.govt.nz>; Aimee Riddell <Aimee.Riddell@regulation.govt.nz>
Cc: Pip Van Der Scheer <Pip.VanDerScheer@regulation.govt.nz>; Laura Fair <Laura.Fair@regulation.govt.nz>; Isabelle Sin <Isabelle.Sin@regulation.govt.nz>
Subject: RE: AI usage by provider

s 9(2)(h)

Cheers,

Jeremy

Jeremy Shoebridge (he/him)
Acting Head of Legal
Ministry for Regulation
īmēra: jeremy.shoebridge@regulation.govt.nz | **waea pukoro:** s 9(2)(a) – call only, not text

www.regulation.govt.nz

From: Adam Jackson <Adam.Jackson@regulation.govt.nz>
Sent: Thursday, 30 January 2025 3:04 pm
To: Aimee Riddell <Aimee.Riddell@regulation.govt.nz>; Jeremy Shoebridge <Jeremy.Shoebridge@regulation.govt.nz>
Cc: Pip Van Der Scheer <Pip.VanDerScheer@regulation.govt.nz>; Laura Fair <Laura.Fair@regulation.govt.nz>; Isabelle Sin <Isabelle.Sin@regulation.govt.nz>
Subject: AI usage by provider

Hi both,

As discussed with Aimee earlier, we're exploring whether our service provider can use AI to help us with submissions analysis.

Aimee thought it sounded fine, so long as we cover off the following questions:

(a) Will the data we provide to our Service Provider (Public Voice) be used to train the model for future use? (No is the answer we want.)
(b) Where are the servers for the AI located? (Hopefully not some dodgy place.)
(c) How will they construct the prompts used to analyse the data? (Ideally suggesting a human-like analysis.)
(d) How will they mitigate the risk of AI error (eg hallucination)?

s 9(2)(h) If not, I'll send the provider an email inquiring about these things shortly – other than the first one, for which they have already told us the answer is "no" and confirmed in writing. (They've also confirmed that the model has SSL.TLS 1.3 encryption for data transmission and zero data retention.)

The answer to (c) will almost certainly be that they are constructed using the tags that we (humans) provide, which Aimee is happy with.

Andrew is on board with AI use, subject to the above, and we know the Minister supports us using it.

Thanks
Adam

Adam Jackson
Chief Advisor to the Chief Executive
Ministry for Regulation
Mobile: s 9(2)(a) **| Email:** adam.jackson@regulation.govt.nz

Ministry for Regulation
Te Manatū Waeture
www.regulation.govt.nz

From: Aimee Riddell
Sent: Friday, 31 January 2025 1:46 pm
To: Adam Jackson
Cc: Isabelle Sin; Pip Van Der Scheer; Laura Fair
Subject: RE: RSB Submission Analysis

Thanks Adam, also ok with these responses. Endorsing AI usage as per MfR requirements.

Cheers,
Aimee

From: Adam Jackson <Adam.Jackson@regulation.govt.nz>
Sent: Friday, 31 January 2025 1:40 pm
To: Aimee Riddell <Aimee.Riddell@regulation.govt.nz>
Cc: Isabelle Sin <Isabelle.Sin@regulation.govt.nz>; Pip Van Der Scheer <Pip.VanDerScheer@regulation.govt.nz>; Laura Fair <Laura.Fair@regulation.govt.nz>
Subject: FW: RSB Submission Analysis

Hi Aimee,

I'm happy with these responses. Are you? If so, I think we have a green light on the AI usage.

Adam Jackson
Chief Advisor to the Chief Executive
Ministry for Regulation
Mobile: s 9(2)(a) **| Email:** adam.jackson@regulation.govt.nz

Ministry for Regulation
Te Manatū Waeture
www.regulation.govt.nz

From: s 9(2)(a) @publicvoice.co.nz>
Sent: Friday, 31 January 2025 11:34 am
To: Adam Jackson <Adam.Jackson@regulation.govt.nz>
Subject: RE: RSB Submission Analysis

Hi Adam,

Here are our responses to your questions about information handling and AI use:

(a) Confirmed - information will not be used to train AI models.
(b) Confirmed - all information will be destroyed upon project completion.
(c) Data processing will occur in New Zealand, European Union or United States.
(d) Maximum of 3 PublicVoice staff members will have access to the information.
(e) All staff accessing the data will be based in New Zealand. Zero data retention policies ensure only PublicVoice staff have access.

(f) Our approach to AI prompts will be:

- Converting your tag framework into structured analysis prompts
- Using standardised templates to ensure consistent analysis
- Testing prompts extensively before full implementation
- Regular quality checks of AI outputs against human validation
- No sharing of prompts or methodology outside the project team

(g) To mitigate AI error risks:

- AI outputs will be human-reviewed by our analysts
- Implementing multi-step validation process
- Using MAXQDA's built teamwork tools
- Regular cross-checking between team members
- Maintaining audit trail of all analysis steps
- Immediate flagging and correction of any inconsistencies

(h) Confirmed - all information will be kept strictly confidential in accordance with our privacy and confidentiality policies. See Privacy Policy | PublicVoice Research & Consultation NZ

Best regards,
s 9(2)(a)

From: Adam Jackson <Adam.Jackson@regulation.govt.nz>
Sent: Friday, 31 January 2025 8:40 am
To: s 9(2)(a) @publicvoice.co.nz>
Subject: RE: RSB Submission Analysis

Thanks s 9(2)(a) . s 9(2)(g)(i)

Adam Jackson
Chief Advisor to the Chief Executive
Ministry for Regulation
Mobile: s 9(2)(a) **| Email:** adam.jackson@regulation.govt.nz

Ministry for Regulation
Te Manatū Waeture

www.regulation.govt.nz

From: s 9(2)(a) @publicvoice.co.nz>
Sent: Friday, 31 January 2025 8:38 am
To: Adam Jackson <Adam.Jackson@regulation.govt.nz>
Cc: Isabelle Sin <Isabelle.Sin@regulation.govt.nz>; Laura Fair <Laura.Fair@regulation.govt.nz>; Pip Van Der Scheer <Pip.VanDerScheer@regulation.govt.nz>; Aimee Riddell <Aimee.Riddell@regulation.govt.nz>
Subject: RE: RSB Submission Analysis

Good morning, Adam,

Thanks for sending through these questions. I'll get back to you early this afternoon with responses to these AI and privacy questions, along with the other information we discussed around approach and pricing.

Best,
s 9(2)(a)

From: Adam Jackson <Adam.Jackson@regulation.govt.nz>
Sent: Thursday, 30 January 2025 4:58 pm
To: s 9(2)(a) @publicvoice.co.nz>
Cc: Isabelle Sin <Isabelle.Sin@regulation.govt.nz>; Laura Fair <Laura.Fair@regulation.govt.nz>; Pip Van Der Scheer <Pip.VanDerScheer@regulation.govt.nz>; Aimee Riddell <Aimee.Riddell@regulation.govt.nz>
Subject: RE: RSB Submission Analysis

Hi s 9(2)(a) ,

I've just been working through a few issues around the use of AI with our IT and legal gurus and have set out a few questions/requests for you below. I've also covered off a couple of more general privacy issues. I have a pretty good idea of your answers to most from our conversations and correspondence so far, but just need to check and get your replies in writing for assurance.

(a) Please confirm that the information we provide will not be used to train your AI model. (I know you've already answered this, but if you could reply "confirmed" so we have it all in one place, that would be helpful.)
(b) Please confirm that the information will be destroyed when the job is finished. (Again, I know you've already confirmed ....)
(c) Where will the information be held/processed by the AI – as in which country(ies)? This question is just to make sure that we don't have personal information, etc, held in countries with poor data protections.
(d) How many people are likely to have access the information? (A broader privacy question, rather than an AI one.)
(e) In which country(ies) will people who have access to the information be based? (Again, a broader privacy question rather than an AI one.)
(f) How will you construct the prompts to get the AI to analyse the data? (For example, will you be taking our tags and turning them into prompts?)
(g) Can you give us a brief description of how you will mitigate the risk of AI error, including by hallucination?
(h) Can you please confirm that you will keep the information we provide to you confidential, for privacy and general confidentiality reasons?

Based on what you've told us already, I don't anticipate your answers throwing up any issues, though we haven't previously talked about (c)/(e) – in which case we'll be comfortable with you using AI to help do the analysis.

Kind regards
Adam

Adam Jackson
Chief Advisor to the Chief Executive
Ministry for Regulation
Mobile: s 9(2)(a) **| Email:** adam.jackson@regulation.govt.nz

20250309 Rules for including submissions and other useful stuff

Email submissions

- Email submissions were included only if they were sent within 32 hours of submissions closing or if an extension was explicitly granted.
- Inclusion of submissions:
    - Only one submission was retained from each submitter (based on combination of email address and name).
    - Where instructions were given on which of multiple submissions to include, we followed these.
    - Where multiple submissions were identical or very similar and no instructions were given, we kept the last one.
    - Where multiple submissions differed substantially and no instructions were given, we kept the one with the most content.
    - Where the same submitter submitted by email and CS with different content, we combined the two into one submission.
    - Where submitters sent corrections to earlier submissions, we made these changes.
    - All emails from each email address that sent more than one email were manually examined to implement this process, and emails and submitter names were compared across email submissions and CS submissions.
    - Where a submission was sent in a format we could not access, we invited the submitter to resubmit and included their resubmission even if it was sent after the closing date. Such submissions were identified through an AI procedure combined with selected manual examination.
- We made every reasonable effort to exclude from consideration emails that were not RSB submissions.
    - An AI model was also used to identify potential non-RSB submissions, all of which were then verified manually.
    - Senders who submitted on a different bill/issue were notified and invited to resubmit on the RSB. We dropped such submissions unless the submitter instructed us to retain them and consider them RSB submissions.
- Treatment of multiple submissions
    - Where one email contained multiple separate submissions, we treated each as a separate submission. These were identified primarily by an AI model, and each case was confirmed by a human.
    - Where one submission was on behalf of multiple people/contained multiple signatures, we treated it as a single submission.

Citizen Space submissions

- Where a person filled out the survey more than once (based on name and email), we kept only their last response.
- Submissions where none of the questions were answered were retained.
- Content was considered regardless of which field it was entered in.

Coding of identity

- Identity was a collected field in Citizen Space. Submissions where this field was missing were considered individual submissions if views were expressed in first person singular. In the remaining 24 cases, we assumed the submission was on behalf on an individual, but could not verify this.
- For emails, the process was as follows:
    - AI was used to assign individual, iwi/hapū, or organisation to each submission.
    - The coding of identity was checked by a human in every case where AI identified the submission as not coming from an individual. Corrections were made where necessary.
    - The small number of cases where AI could not give a definitive answer were coded by humans.

Coding of overall stance for emails

- Emails were given preliminary classifications using a large language model (LLM) following the logic of the classification flowchart.

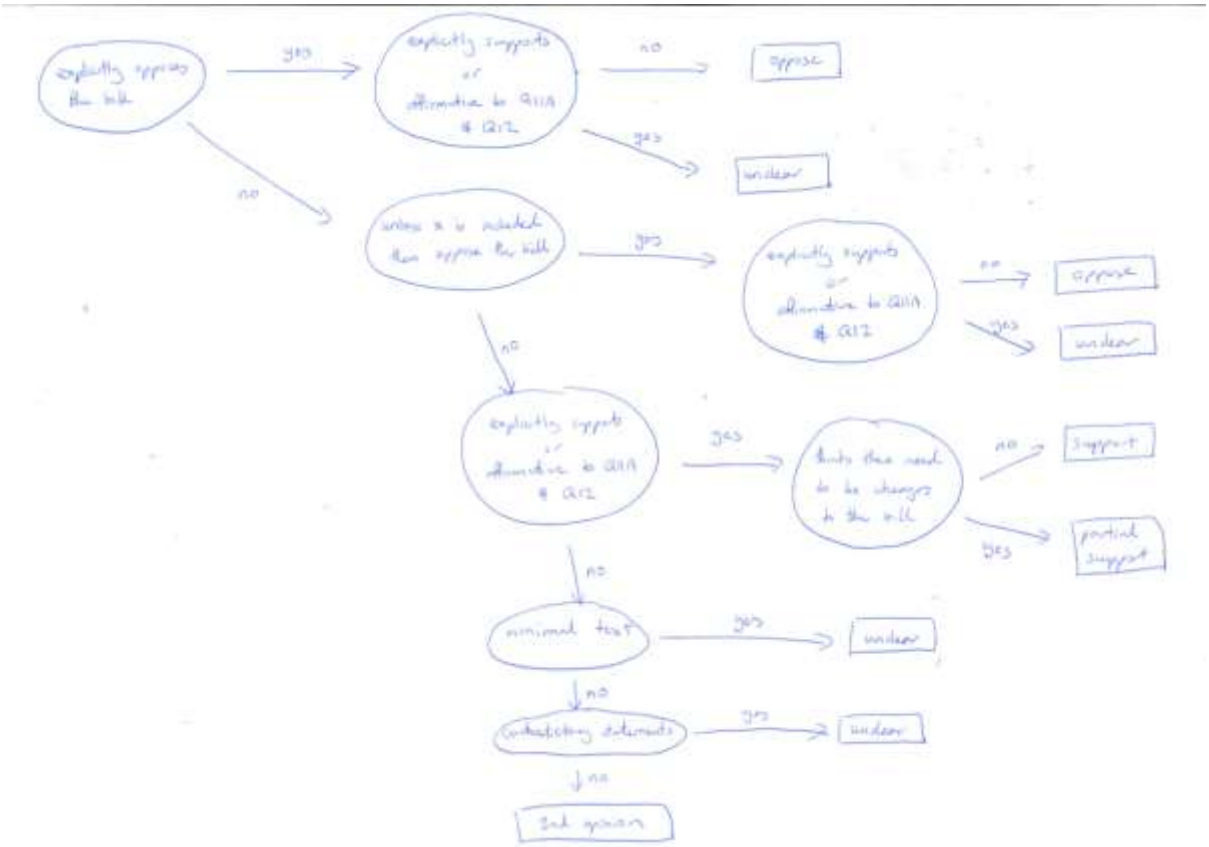

- A random sample of 197 emails (where the probability of inclusion was higher for those not classified as 'oppose') was manually coded.
- The LLM was refined until it produced results that closely matched the manual coding of the 197 emails. Emails the LLM classified as Oppose, Support, or Unclear had a very high probability of being classified the same way manually, but those the LLM classified as partial support were classified manually as partial support only about half the time, and as unclear the rest of the time.
- This version of the LLM was run on the full set of emails.
- Emails classified by the LLM as partial support in the previous step were manually coded.
- Wherever an email had been manually coded, its manual classification was used. In other cases, the LLM classification was used.

Coding of overall stance for CS submissions

- Emails were given preliminary classifications using a large language model (LLM) following the logic of the classification flowchart.
- We manually coded the 567 subs the LLM categorised as something other than oppose plus a sample of 100 the LLM coded as oppose. Coding was done by two separate human taggers, with high levels of agreement.

- Fine tuned the LLM model to achieve a high match rate of tagging with the manually tagged submissions.
- Match between manual and LLM coding was lower for those the LLM coded as partial support or support than for those the LLM coded as oppose or unclear, so we manually coded the rest of the submissions the LLM called partial support or support.
- We modified the AI's coding using several straightforward deterministic rules to increase agreement between the AI's coding and our manual coding.
- Used manual tags for the submissions they were made for, and AI tags for the rest of the submissions.

Substantive submission data set

- Email submissions are included if the length of the email text or length of the attachment text is at least 10,000 characters OR if the submission is on behalf of an iwi/hapū, or organisation.
- CS submissions are included if the combined length of all free text responses plus the length of any associated email submission is at least 10,000 characters OR if the submission is on behalf of an iwi/hapū, or organisation.

Ministry for Regulation
Te Manatū Waeture

Internal policy | Artificial Intelligence

Version	1.0	**Contact**	Digital and Insights Team
Policy Owner	DCE, Organisational Enablement	**Approved**	29th April 2025
SharePoint	Internal policies	**Due for Revision**	April 2026

Context

The Ministry for Regulation (MfR) acknowledges the transformative potential of Artificial Intelligence (AI) in enhancing our operations, boosting efficiency, fostering innovation, and elevating the quality of advice we provide.

We are enthusiastic about embracing these technologies to unlock new opportunities and drive positive change. At the same time, we recognise the inherent risks associated with AI.

This policy is designed to empower MfR staff to responsibly adopt and maximise the benefits of AI, while ensuring its use aligns with principles of safety, transparency, and ethics, and upholds the Ministry's Social License to Operate.

Scope

This policy applies to all MfR **staff** (permanent employees, fixed term employees, secondees, consultants and contractors) at the Ministry for Regulation (**the Ministry** or **we/our**) when using artificial intelligence (AI) to create or process information for the Ministry.

AI is a broad discipline with multiple branches, all focused on creating machines capable of augmenting human intelligence. AI includes Machine Learning (ML), Generative AI (GenAI), Large Language Models (LLM) and Generative Pretrained Transformers (GPT).

The primary focus of ML is to enable machines to learn from past data, improve their performance, and make decisions without explicit coding. Google's search algorithm is an example of ML in its use of past data to refine search results. ML also represents an example of 'narrow AI' which focuses on specific tasks.

GenAI and its subsequent forms, LLM and GPT can process inputs to generate and construct new data. These fall under the category of **'General AI Systems'** which can understand, learn and apply knowledge in multiple domains and can solve problems using machine equivalents of human reasoning, 'common sense', abstract/contextual understanding.

This policy therefore applies to the use and application of all 'General AI Systems' such as Copilot, ChatGPT, Open AI, Gemini, DALL-E and Claude, herein referred to as an AI system.

This policy is also to be considered in conjunction with:

- the requirements of the information and records policy Internal policy | Information and Records Management Policy
- the requirement for acceptable use by staff of Ministry information systems in the acceptable use policy Internal policy | Acceptable Use Policy;
- the information security requirements in the protective security policy Internal policy | Protective Security;
- privacy protection in the privacy policy Internal policy | Privacy.

Principles

Background

With the recent increase in the availability and potential of AI to transform how our business can interact, engage and operate, the Ministry has opportunities to boost productivity, augment staff capabilities, improve the quality of Ministry advice, and to more efficiently and effectively deliver Ministry goals.

As AI systems continue to evolve, developing greater predictive capabilities, there is a need to ensure that AI is utilised in a safe, transparent, ethical, and just way that reflects the Ministry's Social License to Operate (SLO).

There are however risks associated with AI usage which need to be managed to support and empower Ministry for Regulation ('the Ministry') staff to innovate, safely adopt, and derive benefits from using AI systems, including:

- Ensuring Ministry staff act in responsible ways that align with the Ministry's existing policies by setting clear expectations for the use of AI systems,
- Continuing to safeguard the confidentiality, integrity and availability of its information,
- Maintaining the privacy of personal information it holds,
- Ensuring the Ministry retains ownership of and responsibility for its advice,
- Ensuring usage is aligned to the government's Māori Data Governance model: Co-designing Māori data governance - data.govt.nz and
- Ensuring AI results and recommendations are subject to oversight by accountable staff with appropriate authority and capability at every stage.

As per the NZ Information Security Manual (NZISM), the Ministry's Chief Information Security Officer (CISO) is responsible for setting the strategic direction for information security within the Ministry. While some public sector agencies have opted to ban the use of AI systems, the Ministry in consultation with our CISO, has endorsed the use of authorised AI systems on Ministry devices. This is so we don't create stigma or fear in a technology area that is continually evolving.

Principles

The public service System Lead for AI is the Government Chief Digital Officer (GCDO). The Office of the GCDO provides a Public Service AI Framework and guidance for the public service at this link: Public Service AI Framework | NZ Digital government

The OECD's values-based AI principles inform the principles of the Public Service AI Framework:

Inclusive, sustainable development

Public Service AI systems should contribute to inclusive growth and sustainable development through a focus on innovation, efficiency and resilience, and on reducing economic, social, gender and other inequalities and protecting natural environments. AI use should consider and address concerns about unequal access to technology.

Human-centred values

Public Service AI use should respect the rule of law, democratic values and human rights and labour rights through the lifecycle of each AI system or product. These rights and laws include personal data protection and privacy, dignity, non-discrimination and equality, self-determination and autonomy. Public service workers have the right to be consulted on changes made to their work and working arrangements. Agencies need to provide human oversight throughout the AI lifecycle to ensure ethical and appropriate use.

Transparency and explainability

The Public Service needs to commit to transparency in its use of AI. People interacting with government AI systems or receiving AI-assisted services should be aware of and understand how AI is being used. To support this, agencies should publicly disclose when AI systems are used, how they were developed and how they affect outcomes — as relevant and appropriate according to the given use case. Agencies should also enable people affected by the outcome of an AI system to understand how the outcome was determined.

Safety and security

Public Service AI systems should treat the security of customers and staff as a core business requirement, not just a technical feature (security-by-design). They should

minimise risk to individual or national safety and security under normal use, misuse or adverse conditions. The Public Service should ensure traceability of data, apply a robust risk management approach and work collaboratively with commercial and security colleagues in the procurement and assurance of AI tools.

Accountability

AI use within the Public Service should be subject to oversight by accountable humans with appropriate authority and capability at every stage. This should include the application of relevant regulatory and governance frameworks, reporting, auditing and/or independent reviews.

Agency AI capabilities need to keep pace with technological changes, to maintain a strong understanding of AI systems and their limitations.

The Ministry commits to regularly reviewing, refreshing, and re-publishing these guidelines to reflect updated guidance from the Office of the GCDO, updated Ministry policy advice, developments in technology, opportunities and risks.

Implementing this policy

The following expectations are aligned to the Ministry's Internal policy | Information and Records Management Policy, Internal policy | Protective Security, Internal policy | Privacy and Internal policy | Acceptable Use Policy policies.

Provided Ministry staff comply with these guidelines, the risks are acceptable when compared to the benefits that are likely to be gained from responsible use of AI systems.

1. **Use of Ministry devices:** For work purposes, a Ministry-managed device must be used to access only AI systems on the Ministry's Allowlist. Note that the Ministry already blocks access via its IT security firewall to some AI systems (eg DeepSeek) until the completion of a satisfactory cyber security, information and privacy risk assessment.

2. **Classified information:** Official Ministry information, classified, personal or other information that would not normally be publicly available must not be 'fed into', submitted, or provided to, any AI system except for Microsoft Copilot because it operates within the Ministry's protected M365 tenancy. Staff must apply the same security best practices used for all Ministry information and data.

3. **Registration**: The Ministry's staff email address must be used when using AI systems for Ministry business purposes. This enables the Ministry to understand system performance, usage, associated costs and respond to requests for information on Ministry AI usage or any investigative needs.

4. **Protect Māori data sovereignty**: The Ministry has an expectation to act in accordance with Te Tiriti o Waitangi principles. The input and/or production of data and information pertaining to Māori people, language, culture, resources or environments must be done in accordance with the government's Māori Data Governance model and consultation with, or under the advisement of, established Te Tiriti partners to understand and actively manage the impacts of AI for Māori. Ministry staff should be aware that current AI systems may have omissions in authentically representing indigenous cultures. Ministry staff must consult with the Digital and Insights team on appropriate protocols.

5. **Information breach:** Any information breach (or concern that such has occurred) must be reported immediately, in accordance with the Ministry's Internal policy | Protective Security

6. **Decision making:** An AI system must not be empowered to make a business decision.

7. **Use good judgement and validate outputs:** Ministry staff must judge whether the use of an AI system is appropriate, and appropriately scrutinise, validate, and verify any output from an AI system to be used by the Ministry.

8. **Disclosure**: If an AI system has been used to produce a document, then the contribution from the AI system must be disclosed within the document as an integral part of that document's provenance. Identify AI generated text in a footnote in formal documents.

9. **Compliance with security policies:** When using AI systems, Ministry staff must use the same security practices used for all Ministry information. This includes using strong passwords, keeping software up-to-date, and following the Ministry's Internal policy | Protective Security, Internal policy | Information and Records Management Policy and Internal policy | Privacy policies.

10. **Ethical considerations:** In addition to the statements above, use of AI systems must align with the Algorithm charter for Aotearoa New Zealand Charter and be transparent. Ethics and human rights must be considered.

11. **Ministry staff must understand the risks of using AI systems:**
    - AI systems can get things wrong and 'hallucinate' incorrect facts
    - AI systems can be biased and gullible when responding to leading questions
    - The Ministry has an obligation to consider Māori perspectives in our work; AI system bias can raise questions regarding Māori and indigenous information sovereignty, which can breach Māori tikanga by undermining Māori rangatiratanga
    - AI systems can be coaxed into creating toxic content and can be prone to 'injection attacks'
    - AI systems can store all the information submitted to it, including the identity of requestor; and once information is submitted to the AI system, the Ministry can expect to have no control of the information or how it is used
    - AI systems are rapidly evolving, risks can be underexplored, and new developments can bring new risks
    - Public attitudes – including social licence – towards AI in general is very unclear
    - An AI system is only as good as the information upon which it is trained.

12. **Assume human intervention**: Ministry staff must always assume that another human has access to interactions with AI systems. Ministry staff must be mindful of the information provided and how it might be used maliciously to reflect

poorly on yourself, others, or the Ministry. Ministry staff must ensure interactions only contain information that is already publicly discoverable.

13. **Do not use GenAI for legal advice or guidance:** AI systems must not be used to provide legal advice. However, AI systems may be used to help summarise legislation, notes or legal research and commentary.

14. **Automation:** Only Ministry Allowlist AI systems must be used to assist with automation such as handling repetitive tasks, processing or profiling information for contact or customer relationship management, scheduling appointments or processing information. These activities must be managed by Ministry Allowlist AI systems where there are sufficient agreements and information protections in place e.g. Copilot and automation tools such as Power BI and Power Automate available within the Ministry's M365 platform.

15. **New AI systems:** The Ministry is open to critically evaluating any request by Ministry staff who are interested in using a specific AI system for their Ministry activities. Any such request must be made to the Head of Digital and Insights. Following a cyber, information and privacy risk assessment, the AI system will be considered for inclusion on the Allowlist.

Ministry staff can do this...

✅ Copilot is the preferred AI tool for Ministry use. The Ministry has invested and will continue to invest heavily in Microsoft 365 (M365) as its strategic productivity and collaboration platform. Ministry staff can use Copilot to analyse and summarise Ministry information including submissions data because, by operating on data only in the Ministry's M365 tenancy, this means that the safety, security and control of data remains with the Ministry.

✅ Ministry staff can also access ChatGPT, Gemini, Perplexity and other AI systems on our "Allow List" via your browser.

✅ Ministry staff can ask ChatGPT, Gemini, Perplexity and other AI systems on our "Allow List" questions on information already in the public domain.

- eg can you provide me with a Risk Management framework to assess environmental risks?

- eg can you provide a summary, from the New Zealand Ministry for Regulation Strategic Intent 2024/25-2028/29 document published on its website, of its role in regulatory system leadership?
- eg can you give me a template for a Project Brief for initiation?
- eg who are the regulatory agencies in New Zealand?
- eg what is the purpose and objectives of the New Zealand Financial Markets Authority (FMA)?
- eg can you provide a comparison between hairdressing regulations in Sweden and New Zealand?

Except for Copilot, Ministry staff can't do this...

❌ Upload any data in any form to other AI systems eg ChatGPT, Gemini, Perplexity

❌ This means, no we can't ask other AI systems to summarise Ministry information that is not already in the public domain.

❌ So, this means, for example, we can t upload submissions data to other AI systems eg ChatGPT, Gemini, Perplexity.

Related policies and more information

1. Public Service AI Framework | NZ Digital government
2. Co-designing Māori data governance - data.govt.nz
3. Algorithm charter for Aotearoa New Zealand - data.govt.nz
4. Internal policy | Information and Records Management Policy
5. Internal policy | Protective Security
6. Internal policy | Privacy
7. Internal policy | Acceptable Use Policy

Glossary

The following terms are used in this policy:

- **Allowlist**: An allowlist, also known as a whitelist, is a security measure that specifies a list of trusted entities (like IP addresses, websites, applications, or email addresses) that are granted permission to access a system or network, while all others are denied access by default. For the Ministry this means automatically blocking all AI systems by default and then only permitting those we wish to allow.

- **Artificial Intelligence (AI)**: The field of software engineering that creates services that, without explicit programming, can generate outputs for particular sets of inputs.

- **Generative AI (GenAI)**: A system that once prompted or questioned generates text or images or other content that closely resembles human-created content. GenAI works by matching user prompts to patterns selectively downloaded from the Web within a Large Language Model (LLM), then using 'neural networks' to probabilistically fill in the blank', along the lines of predictive text messaging. ChatGPT is an example of a GenAI service. Many other GenAI services are available or are under development.

- **Hallucination**: A response by an AI system that may be false or distorted because of the characteristics of the content within the LLM, or the neural network used within the AI system.

- **Injection attack**: A cyberattack where an attacker supplies untrusted input to an AI system which alters the underlying LLM or the course of the system's execution, and allows attackers to access, steal, or compromise the system's information, the system itself, or users' information.

- **IT Security Firewall**: A network security device that monitors and filters incoming and outgoing network traffic based on an organisation's previously

established security policies. At its most basic, a firewall is essentially the barrier that sits between a private internal network and the public Internet to protect a network or system from unauthorized access.

- **Māori Data Sovereignty**: Refers to the inherent rights and interests of Māori in relation to the collection, ownership and application of Māori data.

- **M365 Tenancy**: A private dedicated, isolated instance of M365 services, like Office 365, Azure, Intune, etc., assigned to a specific organisation, where all data and user accounts are stored securely.

Appendix 1 – Approved AI systems

The Ministry is enthusiastic to empower Ministry staff to innovate, safely adopt and derive benefits from using AI systems. An appropriate cyber, information and privacy risk assessment has been satisfactorily completed on Ministry approved AI systems.

Product	Description	Approval date	Approved by *
THE MINISTRY'S PREFERRED AI TOOL is COPILOT which operates within the Ministry's M365 secure platform			
Microsoft Copilot Studio	Microsoft Copilot Studio is a powerful, cloud-based platform that allows organisations to build, customise, and deploy AI-powered copilots and autonomous agents tailored to their business needs. It's part of the Microsoft Power Platform and integrates deeply with Microsoft 365, Dynamics 365, and other enterprise systems.	23/05/2025	CISO/HoDI

Microsoft Designer	Microsoft Designer is a web-based graphic design tool powered by Copilot AI, designed to help users quickly create professional-quality visuals for social media, presentations, marketing materials, and more—without needing advanced design skills.	23/05/2025	CISO/HoDI
Microsoft 365 Copilot Chat	Copilot Chat is a conversational AI feature within Microsoft 365 that allows users to interact with their work data using natural language. It's part of the broader Microsoft Copilot experience and is designed to help users be more productive by making it easy to ask questions, get summaries, and automate tasks—all through a chat interface. It can access and reason over your emails, documents, meetings, and chats (with appropriate permissions) to provide relevant, personalised responses. Ask things like "What were the key points from last week's meeting?" or "Summarize the latest project update email." Summarise long email threads or documents. Draft emails, reports, or presentations based on your prompts. Analyse Excel spreadsheets and generate insights or visualisations. Create formulas or pivot tables based on natural language queries. Schedule meetings, set reminders, or manage tasks in Outlook and Teams. Help prepare for meetings by summarising past conversations and documents. You can chat with Copilot in a Teams-like interface to ask questions or give instructions. It understands context from your Microsoft 365 environment, making it more personalised and relevant.	23/05/2025	CISO/HoDI
Microsoft 365 Copilot (Paid version)	Microsoft 365 Copilot—is a premium AI assistant designed to enhance productivity, creativity, and decision-making across an organisation. It integrates deeply with Microsoft 365 apps like Word, Excel, Outlook, Teams, and PowerPoint, and is tailored for enterprise environments with robust security, compliance,	23/05/2025	CISO/HoDI

	and management features. It uses natural language to generate content, summarise documents, analyse data, and automate tasks in Word, Excel, Outlook, PowerPoint, and Teams. It has a conversational interface that allows users to interact with their work data and documents using AI-powered chat. You can build and manage custom AI agents tailored to your business needs, including SharePoint-based agents and integrations via Microsoft Graph connectors. It has built-in data protection, IT management controls, and compliance with Microsoft's enterprise security standards. Copilot reasons over personal work data (emails, files, meetings) to provide context-aware assistance.		
STAFF CAN USE THESE TOOLS BELOW, BUT MUST **NOT** UPLOAD MINISTRY DATA INTO THE AI SYSTEMS			
Adobe Express	Adobe Express is a user friendly, web-based content creation platform designed for anyone —from beginners to professionals—who wants to quickly create high-quality graphics, videos, and documents. It's especially popular among social media marketers, educators, small businesses, and content creators.	23/05/2025	CISO/HoDI
Adobe Sensei	Adobe Sensei is Adobe's AI engine that powers smart features across its apps to help users create, edit, and analyse content more quickly and intelligently. It is generally built into other Adobe products such as photoshop and lightroom rather than being an app in itself.	23/05/2025	CISO/HoDI
ChatGPT	ChatGPT is based on a large language model (like GPT-4), trained on vast amounts of text data. It doesn't "know" things like a human does, but it can generate highly relevant and coherent responses based on patterns in the data it was trained on.	23/05/2025	CISO/HoDI

Anthropic Claude	Anthropic Claude is a family of advanced AI models developed by Anthropic, a company founded by former OpenAI researchers. As of 2025, the latest generation is the Claude 4 series, which includes two models: Claude Opus 4 and Claude Sonnet 4. These models are designed to be powerful, safe, and capable of handling complex, long-running tasks with minimal human input. It can: Understand and generate human-like text, perform deep reasoning and analysis, automate workflows and long-term tasks, write and debug code, summarise, search, and synthesise large volumes of information.	23/05/2025	CISO/HoDI
Google Gemini	Gemini (formerly Google Bard) is Google's conversational AI chatbot. Gemini runs on Google's family of multimodal AI models to understand and generate text, and work across other mediums like images audio, and video.	23/05/2025	CISO/HoDI
Perplexity	Perplexity AI is an AI-driven search engine and chatbot that uses large language models (LLMs) to answer user queries by drawing information from the web and providing cited sources within its responses. The AI model combines a traditional search engine with an AI assistant, delivering answers in natural language backed by references.	23/05/2025	CISO/HoDI
Napkin AI	Napkin AI is a creative and organisational tool designed to help users capture, connect, and reflect on their ideas using artificial intelligence. It acts like a personal thinking partner, ideal for writers, researchers, creatives, and anyone who wants to make sense of scattered thoughts or inspirations.	23/05/2025	CISO/HoDI

*** CISO/HoDI = Chief Information Security Officer | Head of Data & Insights**