14 July 2025

s 9(2)(a)

Official information request
Our ref: R00979

Tēnā koe s 9(2)(a)

Thank you for your Official Information Act 1982 (OIA) request received on 22 May 2025. You requested information pertaining to the methodology and techniques used to analyse submissions on the discussion document on the proposed Regulatory Standards Bill (RSB). Your full request is outlined in **Appendix A**, but includes:

- *the generative model and various prompts used to analyse submissions.*
- *what AI tools were used,*
- *what prompts were used when analysing submissions,*
- *why AI was chosen to analyse these submissions,*
- *who gave the approval for AI to be used to analyse the submissions,*
- *whether it is Ministry policy to use AI in this fashion.*
- *how often AI is used within the Ministry of Regulation and why?*
- *what costs are associated with these tools,*
- *what security measures are in place to ensure private government information is not accessed, used, or retained by the AI tools?*
- *what your measures for success when using generative AI are,*
- *the expected outcomes,*
- *how you verify receiving accurate data?*

Tools and prompts used to analyse submissions

Details about the methodology and techniques used to analyse submissions are available in the summary of submissions document which is available on the Ministry for Regulation

(the Ministry) website[1] (page 42). Your request for this information is refused under section 18(d) of the OIA, as the information is publicly available.

Prompts used in the analysis of submissions on the discussion document for the proposed Regulatory Standards Bill are detailed in **Appendix B**.

Use of AI for submissions on discussion document

Use of AI for analysis of submissions has been approved by the Ministry's Chief Information Security Officer/Head of Digital and Insights.

The Ministry worked with Public Voice to ensure the privacy and security of submissions. The following privacy and security measures were implemented by Public Voice:

- Inputs were not used to train other Large Language Models (LLM) models
- All information was kept strictly confidential in accordance with privacy and confidentiality policies, and destroyed upon project completion
- A maximum of three staff members had access to the information, all based in New Zealand.

Our objective was to analyse all submissions in such a way that each unique point of view was identified, that the level of support for each point of view was measured, and that these could then be considered in the policy process.

Tools like software, machine learning and AI help us to do this far more accurately and efficiently than ever before – by doing sorting, grouping and measuring tasks at scale and with more speed and accuracy than humans. This frees up human time to do policy analysis and provides more thorough, detailed and accurate analytics to inform that policy analysis.

One example of time saving relates to submissions that feature similar or identical wording that has been provided by interested parties to use in individual submissions. Such submissions are certainly valid and need to be counted along with other submissions. However, they can be identified and counted using software tools. Having a human re-read each instance of this repetition is not necessary and nor is it often a good use of time.

We appreciate this is a new use of technology in our democratic process and that it may be uncomfortable for some people. Most, if not all submissions were also viewed by a human analyst wherever possible, to ensure that the insights from the technology were accurate.

[1] https://www.regulation.govt.nz/assets/Publication-Documents/Information-Release-Summary-of-Submissions-for-proposed-Regulatory-Standards-Bill.pdf

General use of AI within the Ministry for Regulation

The Ministry's AI policy, which I have enclosed as **Appendix C**, 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. However, the Ministry for Regulation does not keep a register of how often AI is used by its staff, or for what purpose. Your request for this information is refused under section 18(g) of the OIA as the information is not held by the Ministry for Regulation and we do not believe it is held by another agency.

The Ministry is trialling and using AI technology for everyday work. This is in line with the Government Chief Digital Officer's *Public Service AI Framework*, which sets expectations for the safe, ethical, and effective use of AI across the public sector and the *Guidance for the Safe Use of AI in the Public Sector*. Use cases have include reviewing research literature and draft documents, thematic research and document analysis, transcribing internal learning videos and document comparison and summarisation.

The main AI tool currently used, Microsoft 365 Copilot Chat, is provided through the Ministry's M365 desktop and collaboration software environment at no additional cost. The cost of an upgraded Microsoft 365 Copilot subscription for one staff member is $567.60 per year.

The Ministry's information and analysis results are not used to train AI and is stored in approved, resilient and secure information systems within its private storage facilities (tenancy) on an enterprise-strength platform. Robust security protections, monitoring and alerting have been applied to ensure compliance with GCSB NZISM standards. Information retained within an AI tool such as Copilot are held within the Ministry's tenancy and are not available to external parties.

Additionally, the Ministry refers to the all-of-Government guidelines on the use of technological tools (including artificial intelligence and LLMs), including:

- Responsible AI Guidance for the Public Service (digital.govt.nz)[2]
- The Artificial Intelligence Guidance (data.govt.nz)[3]
- Artificial intelligence and the Information Privacy Principles[4]
- 2024 cross-agency survey of use cases for artificial intelligence (digital.govt.nz)[5]
- Public Service AI framework (digital.govt.nz)[6].

[2] https://www.digital.govt.nz/assets/Standards-guidance/Technology-and-architecture/Generative-AI/Responsible-AI-Guidance-for-the-Public-Service-GenAI-Print.pdf
[3] https://data.govt.nz/leadership/centre-for-data-ethics-and-innovation/guidance/artificial-intelligence-guidance
[4] https://www.privacy.org.nz/assets/New-order/Resources-/Publications/Guidance-resources/AI-Guidance-Resources-/AI-and-the-Information-Privacy-Principles.pdf
[5] https://www.digital.govt.nz/dmsdocument/262~full-results-2024-cross-agency-survey-for-artificial-intelligence-ai-use-cases/html
[6] https://www.digital.govt.nz/standards-and-guidance/technology-and-architecture/artificial-intelligence/public-service-artificial-intelligence-framework

Measures, outcomes and verification of AI information

The Ministry's measures for success for the use of AI include AI Policy alignment with principles of safety, transparency, and ethics, and upholds the Ministry's social license to operate, technical performance and user engagement, operational efficiency and business impact.

This technology has the potential to significantly accelerate our analysis of complex regulatory systems. Use cases to date have delivered time savings through process efficiencies, quicker and broader thematic analysis, data product visualisations, quicker legislation analysis and comparison. As previously mentioned, by automating time-consuming tasks and surfacing insights faster, we can enable our teams to focus their time and skill on higher-value work.

Regarding the verification of information, the Ministry's AI Policy makes it clear that staff should not rely solely on AI output: '*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*'.

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

s 9(2)(a)

22/05/2025

OIA Requests Team

Ministry for Regulation

PO Box 577

Wellington 6140

hello@regulation.govt.nz

Dear Sir or Madam

Official information request: Elaboration of the use of large language models (generative AI) with regards to the Regulatory Standards Bill

Please supply the following information under the Official information Act (OIA):

It has come to my attention that the Ministry of Regulation has used generative AI to analyse a majority of the 23,000 submissions it received on the Regulatory Standards Bill consultation document.

I am formally requesting information pertaining to the methodology and techniques used with the large language model to analyse these submissions, including the generative model and various prompts used to analyse submissions.

Could you please inform me what AI tools were used, what prompts were used when analysing submissions, why AI was chosen to analyse these submissions, who gave the approval for AI to be used to analyse the submissions, and whether it is Ministry policy to use AI in this fashion.

Could you please let me know how often AI is used within the Ministry of Regulation and why? Additionally, what costs are associated with these tools, and what security measures are in place to ensure private government information is not accessed, used, or retained by the AI tools?

Finally, could you please let me know what your measures for success when using generative AI are, the expected outcomes, and how you verify receiving accurate data?

Yours faithfully

s 9(2)(a)

Appendix B

Prompts used for submissions received via email

Classify each submission into exactly one of four categories:
1. Support
2. Partial support
3. Oppose
4. Unclear

IMPORTANT CLASSIFICATION CRITERIA:

SUPPORT

- The submission explicitly endorses or strongly praises the Bill as a whole.
- Uses phrases like "strongly support," "fully support," "want this Bill passed."
- May express minor concerns but clearly indicates wanting the Bill to pass.
- If the concerns mentioned are actually changes they want before they would be happy with the Bill, classify as Partial_support instead.

PARTIAL_SUPPORT

- Expresses support but with conditions or qualifications of any significance.
- Uses phrases like "support if changes made," "agree with principles but needs modifications."
- Includes cases where they see good aspects of the Bill but have reservations.
- Key indicator: They want the Bill to pass, but only after specific concerns are addressed.

OPPOSE

- Explicitly wants the Bill withdrawn or rejected.
- Uses phrases like "do not want this Bill," "it's harmful," "want it stopped/scrapped," "I oppose."
- States current arrangements are better or sufficient.
- Explicitly aligns with known Bill opponents or supports critics' views.
- If they state "unless X changes, I don't want the Bill to pass," classify as Oppose.
*- **Additionally, classify as Oppose in any of these situations**:*
- If a submission contains explicit rejections of key components of the Bill.
- If a submission has few responses but explicitly rejects a core component of the Bill.

- If a submission expresses satisfaction with current regulation and rejects additional regulatory measures.
- If a submission explicitly states "I do not support" and provides negative responses for key oversight measures.
- If a submission criticises key mechanisms (e.g., consistency checks) as enabling corporate exploitation that harms the environment.
- If a submission criticises the Bill for prioritising economic productivity over social, environmental, and community well-being and calls for alternative principles.
- If a submission explicitly expresses support for the submissions of Jane Kelsey, Jonathan Boston, or the New Zealand Public Service Association Te Pūkenga Here Tikanga Mahi (the PSA).
- If a submission criticises the Bill for prioritising individual property rights and economic productivity over environmental protection and Māori rights.
- If a submission praises existing regulation but rejects proposed changes and insists on keeping the current system unchanged.
- If a submission includes explicit statements such as "I do not support" in key final questions.
- If a submission provides detailed criticisms (e.g., lack of consultation, inadequate environmental protection, undermining democratic processes), even if most responses are NA.
- If a submission includes explicit negative responses (e.g., "No" to oversight measures) or a clear statement of non-support due to dishonoring Te Tiriti o Waitangi.
- If a submission explicitly decries the Bill as a tool for a hidden political or corporate agenda that undermines the rights of citizens and Māori.
- If a submission includes explicit statements of non-support (e.g., "I do not support") in key sections.
- If a submission includes an explicit statement of non-support, such as "I do not support because It dishonours Te Tiriti o Waitangi."
- If a submission contains explicit criticism of the proposed principles as unbalanced or anti-democratic.
- If a submission explicitly states "I do not support it" and highlights concerns like lack of public consultation.

UNCLEAR

- The submission does not clearly indicate support, partial support, or opposition.
- The answers are contradictory or too vague to determine a position.
- Note: If they disagree with aspects of the Bill but don't explicitly oppose it entirely, classify as Unclear.

Prompts used for submissions received via Citizen Space

System Context:

You are an expert in analysing government consultation submissions in New Zealand.

Key points to incorporate:

1. *If a submitter states "regulation is fine, nothing needs to change" (or words to that effect) with no other commentary, classify as "Oppose."*
2. *Dissatisfaction with the current regulatory arrangements alone does NOT necessarily indicate support for the Bill. Some submitters may be referring to the Bill itself when talking about "current regulation." Carefully evaluate their statements to determine if they are indeed supportive, opposed, or uncertain about the new Bill.*
3. *If Q33 = "Yes", default to "Support" unless the submission explicitly indicates they oppose.*
4. *If Q33 = "Yes" but the submission also conveys that the current system is already good (or doesn't really need change), classify as "Partial Support" rather than fully "Support."*
5. *If Q33 = "No", default to "Oppose" unless the submission explicitly indicates they support.*
6. *If Q33 = "Not Answered", classify based on the content of the submission alone; if no clear stance emerges, default to "Unclear".*
7. *If uncertain, default to "Unclear".*

Classification Values:

- *Support, Oppose, Partial Support, Neutral, Unclear*

Key Guidelines:

- ***Support:*** *Explicitly states support for the bill, or Q33 = "Yes" without any contradictory statements*
- ***Oppose:*** *Explicitly states opposition to the bill, or Q33 = "No" without explicit support, or states the status quo is fine with no changes needed*
- ***Partial Support:*** *Q33 = "Yes" but submission text also expresses that current system is sufficient, or they strongly prefer minimal/no legislative changes*
- ***Neutral:*** *Provides feedback but doesn't lean clearly for or against*
- ***Unclear:*** *No explicit stance, ambiguous position, or incomplete response*

Appendix C

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

RELEASED UNDER THE OFFICIAL INFORMATION ACT

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.

RELEASED UNDER THE OFFICIAL INFORMATION ACT 9

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

RELEASED UNDER THE OFFICIAL INFORMATION ACT 9

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

RELEASED UNDER THE OFFICIAL INFORMATION ACT 9

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.

RELEASED UNDER THE OFFICIAL INFORMATION ACT 9

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.

RELEASED UNDER OFFICIAL INFORMATION ACT 9

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

RELEASED UNDER THE OFFICIAL INFORMATION ACT

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.

RELEASED UNDER THE OFFICIAL INFORMATION ACT

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

RELEASED UNDER THE OFFICIAL INFORMATION ACT
9

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

RELEASED UNDER OFFICIAL INFORMATION ACT 9

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.

RELEASED UNDER T
OFFICIAL INFORMATION AC
9

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

RELEASED UNDER T
OFFICIAL INFORMATION AC
9

	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

RELEASED UNDER T
OFFICIAL INFORMATION AC
9

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**

RELEASED UNDER THE OFFICIAL INFORMATION ACT