



# Are AI chats subject to the Official Information Act?



Jodie O'Neill, author

Of course not. Wait... maybe, oh actually definitely....

The question first occurred to me when reading [this Official Information Act request](#) from [Josh Drummond](#)<sup>1</sup> to the Ministry of Education about their use of AI in developing the new education curriculum.

Following an initial information act request which confirmed the Ministry used Copilot in developing the curriculum Drummond had requested *"all tt chat logs relating to curriculum development, the review process that information has undergone to ensure accuracy, reliability and safety and how research or other material derived from or filtered through AI has been utilised in the draft curriculum – including copy+pasting, paraphrasing or other methods."*

As someone who has worked in public

sector adjacent roles for over a decade, I like to think that I'm passably familiar with the Official Information Act. But my initial reaction upon reading this request was that, much like requesting an individual officials' internet search history, chat logs were unlikely to be released under an OIA.

Then I thought about it some more, went back to read the legislation, looked to see if the Ombudsman had released any advice on this issue (not to date) and asked the brains trust of my colleagues. The consensus was in principle yes. In practice, I'm not so sure.

<sup>1</sup> Josh Drummond is a journalist who has written for a number of publications including the Spinoff.



## What's the issue exactly?

To be clear I'm not saying that the Official Information Act doesn't apply to chat logs. Section 2 makes it clear that the Act applies to all information 'held' by an agency. The Ombudsman has clarified that information is 'held' when the information is in an agency's control. There isn't much wiggle room here: if a Government Department has an enterprise version of Copilot, all the chat logs are information held by the agency.

The issue arises in relation to the exemptions available to agencies under the Act. Agencies can refuse to provide information for a range of reasons including on the basis that the request cannot be met without substantial collation or research (section 18(f)) or that it breaches the privacy of an individual (section 9). Requests for internet search histories would normally fall into one or both of these categories.

Upon revisiting Drummond's OIA to write this piece, I was unsurprised to see that the Ministry has responded citing the request is very broad in scope and has asked that the request be refined. This is a pretty standard response to receive from an agency when you request a large amount of information. Drummond refined the request to include any chat prompts relating to a one-month period. Despite this the Ministry declined his request on the grounds that it would require substantial collation. But should that exemption **really be** available for AI chat logs simply because there are a lot of them?

## How did we get here?

The Official Information Act was introduced in 1982 when records were paper based and often poorly indexed. As information moved into electronic systems, collation issues didn't disappear, they changed shape. Volume increased, systems multiplied and information became more fragmented.

It has taken a shockingly long time for paper to be properly excised from our public systems. Even with fully electronic records, the substantial collation exemption is still justified in many cases. Responding to a request doesn't just involve retrieval of the information, but also review and interpretation, considering whether any exemptions apply and redacting information. This can all be extremely resource intensive.

But our systems haven't kept up. If you request information that may be contained in emails, most agencies still rely on the old tried and true method of asking staff to search their inboxes using keywords. While manual and extremely annoying, there are natural limitations (people can only realistically send so many emails in a workday). AI however, introduces a whole world of new challenges.

## Why AI changes the problem

The 2025 Public Service Census reported that only 5% of public servants are using AI in their daily work.<sup>2</sup> But the Government hopes to increase these numbers: New Zealand's Strategy for Artificial Intelligence<sup>3</sup> explicitly calls for public-sector leadership in AI use.

The potential volume of AI chat logs which could be subject to OIA in 5 years alone may be staggering. With the exception of friendly platitudes between colleagues at Christmas and basic administration - most government email correspondence generally contain at least some information that might be of public interest. The same is not true of AI.

I use AI almost daily in my work life and in a typical day I would generate dozens if not hundreds of pages of material. A considerable amount of that would be very low value material – prompts that AI has misunderstood, questions about how to fix my technology problems and repeated prompts or rephrasing which might generate pages of irrelevant content. Much of it would be of little or no public interest.

---

2 [Summary reports and technical information - Te Kawa Mataaho Public Service Commission](#)

3 [New Zealand's strategy for artificial intelligence: Investing with confidence - one page](#)

Given this it seems likely that many agencies will be tempted to reach quickly for section 18(f). But in this day and age, volume alone feels like a weak justification for denying to release information. If AI is being used to produce material that informs public work and decision-making, then transparency demands a better answer than 'there's simply too much of it.'

## What does AI transparency require?

I'm a strong advocate for AI used properly. Knowing whether Government is using AI properly is going to take more than just reading their AI policies. It means a high degree of human-oversight and transparency. It means:

- Understanding the tools being used, their system settings, their strengths, and crucially – their weaknesses.
- Scrutinising the instructions and inputs that are being put into the model.
- Scrutinising the outputs from the model.
- Transparency about the levels of human oversight and checks in place to decide how to use outputs in work products and decision-making.



Put simply, if we can't see the prompts and responses, questions and answers, follow-ups and challenges, then we can't meaningfully scrutinise how AI is being used. This is particularly important at the moment, given the relative immaturity of AI capability across much of the public sector and the lack of developed accountability mechanisms. If agencies are regularly relying on section 18(f) exemptions to respond to requests about chat logs, these risks undermine the intention of the Official Information regime.


## So what can agencies practically do to respond to these challenges?

1. Communicate clear expectations about AI chats being official information and how they should be managed.
2. Explore AI powered retrieval and search capabilities to meet the challenge.
3. *Trial how AI can support analysis, collation and redaction of information.*

The first step is a basic one, but it matters: **agencies need to be clear with their people that AI chat logs in enterprise tools are official information.** If staff are using Copilot at work, their chats are not informal notes or a private working space. They are information held by that agency and sit squarely within the scope of the Official Information Act.

Maybe staff already know this – but I'm willing to bet many don't. The key is not to discourage AI use but to ensure that it's subject to good information management principles from the beginning.

As AI becomes more mainstream, agencies need clearer expectations about how it is used and how outputs are relied upon. Prompts and responses form part of the record of how work is done. Agencies that are thoughtful about this now will be in a far stronger position and less exposed to risks that are coming, if not already here.



The next issue is retrieval. **Most agencies already have access to tools that could support OIA responses involving AI, but they are currently under-utilised and poorly understood.** For those using Copilot, Microsoft Purview Audit logs Copilot activity at an enterprise level. It records who used vvv, when they used it, a summary of the context, and any files or sites Copilot accessed to generate a response. It provides enough detail to scope a request and identify where relevant information sits, but doesn't capture the chat content itself.

However, by default, Purview Audit logs are only retained for 180 days. In my view, that is unlikely to be sufficient to really support agencies to respond to Official Information Act requests. Agencies need to consider whether this is sufficient, or whether they should be investigating longer term solutions.

But finding relevant chats is only half the job. Agencies also need a practical way to access them. It's important to understand that AI chats are not designed for manual retrieval at scale. Most Copilot users' profiles do not have the ability to export their full chat history. The alternative – asking staff to manually copy and paste individual chats is slow, inconsistent and unlikely to scale.

Tools like Microsoft eDiscovery offer a far more practical option, allowing centralised search and retrieval of Copilot chats and relevant material. Crucially, eDiscovery has features allowing for retrieval of full chat content, and associated files. Without tools like this, agencies will increasingly struggle to respond to large AI related requests.

Even with retrieval sorted, volume still remains an issue. As I mentioned earlier, there is still likely to be a substantial volume of material including a significant amount of low-value material. But AI can help to manage this. It can be used to remove obviously irrelevant material and to flag content that might require a closer review, such as personal or commercially sensitive information. This makes it a much more manageable task for a human to review and make final decisions about what should be released, and what should be withheld or redacted.

There's a real opportunity here to completely rethink the way agencies respond to all OIA requests. But so far, we haven't seen any agency take up the opportunity of AI for efficiencies in OIA responses.

## Transparency starts at the prompt

If government is relying on AI to support public work, then transparency requires more than high level assurances that appropriate checks are in place. The Official Information Act should support this by giving access to prompts, and promoting transparency about the outputs and insight into the human judgement that has been applied along the way. Section 18(f) may have been a workable excuse in the age of paper files and fragmented email systems, but it is a weak answer to AI at scale.

But this doesn't have to be a disaster for agencies. The tools already exist to retrieve, review and manage AI generated information. There will undoubtedly be some challenges but there are also so many opportunities.

*Agencies need to prepare now – to test how they can use AI to make their lives easier now, while it is still emerging, rather than waiting until volume, scrutiny and expectations overtake them.*

Get in touch:

**Jodie O'Neill**

[joneill@allenandclarke.com](mailto:joneill@allenandclarke.com)

+64 21 217 6998