

Opinion: Chat and Instant Messages: Rubbish Records on the Rise

Michael Upton



Chat in public sector workplaces has substantially shifted from desktalk conversation to something written down. An activity that did not create records now does, almost completely unintentionally. Of course, people still talk to their workmates, but in many public offices there is now a large amount of typing back and forth every day, peppered with links, jokes, emojis, animated GIFs, and references.

Instant messaging tools have been available on many workers' computers for years, but over the last few years there has been a great wave of increased use. The reasons are obvious: the massive increase in remote working and other flexible arrangements,

with the ubiquity of Microsoft Teams a mix of enabler and catalyst.

When thinking about the benefits of chat you probably first think of working remotely or across geographic locations, but people are finding the same benefits apply when, for example, someone is in a meeting, away sick, or even on the other side of the floor from you. Sending a quick message to a colleague while you remember is often useful, "side chat" in meetings is often seen as less disruptive than talking over someone, and, even if most of a team is collocated, group chats can help to include those who are not present for whatever reason.

My aim with this article is to explore chat from a records management perspective and to argue that the best available option is to automate the deletion of chat after a fixed period of time. I hope this helps others to firm up their thoughts on how they want to manage these records – whether that is because they agree or disagree with what I am about to write.

I will use "chat" to refer to sending and receiving messages that the recipient can see as good as instantly via software. Both Microsoft Teams and Zoom allow for chat in this sense, but there is nothing in here that, in principle, could not be applied to any other application that serves a similar purpose. Many of the qualities of chat are also shared with emails, but chat is much more resistant to any kind of conventional records management. Microsoft Teams has a concept of channel posts and replies, which I still consider chat for the purpose of this opinion piece. There are some differences which I have chosen to put aside for the sake of keeping this brief.

I will refer to public offices and public records, but I expect anything I write will be equally applicable to local authorities and their records.

Chat messages are subject to the Public Records Act 2005, Official Information Act 1982, and Privacy Act 2020

There is no practical reason to debate that chat messages created by a public office are public records. We can just point to the fact that in 2013 the Chief Archivist saw the need to authorise the disposal of “Casual or instantaneous communications that contain information of an ephemeral, non-transactional nature” through General Disposal Authority 7 (“GDA7”).¹ This class of records so neatly describes the vast majority of chat it sounds like fortune-telling. If such communications were not public records, there would be no need to authorise their disposal.

Where GDA7 may not apply, it would have to be because the content of the chat captures something more significant, not less, about the affairs of the public office: a not-so-ephemeral record of a decision, an action taken, or authority given. If the chat is the master (or only) record, that needs to be maintained and kept accessible, for all the usual good reasons.

Chat can also unequivocally meet the broad definitions of “official information” and “personal information” in the Official Information Act and Privacy Act respectively. Screenshots of Microsoft Teams chats are already showing up in published responses to Official Information Act requests. It is commonly understood that New Zealand’s legislation related to information is not limited by the format of the information.

Unintentional, poor-quality records

There are several qualities of chat that make for poor-quality records, even if we decide some of the content is worth capturing and managing.

A key contributing factor to the poor quality is that almost all chat can be described as unintentional records. The records are not the planned product of a recordkeeping choice, but a byproduct of the combination of the relevant regulatory settings and the medium being used. For example, if a person chooses chat to ask a colleague how to use the printer or what time they are finishing work, it is incredibly unlikely that they made that choice because they wanted to keep a good record.

There certainly will be some occasions when people use chat because they see the value of putting something in writing – or, in other words, making a record. These are usually of short-term value.

So, while chat results in records, it is not designed with good recordkeeping or records management in mind. This is true of the technology and of how people use it.

Here are four factors that would make it hard to identify and manage higher-value chat messages.

Information is spread across many separate messages

There has been an ongoing trend in written communication towards sending smaller, more numerous messages. In chat, it would not be unusual for someone to convey one idea through a series of five to ten messages, with a sentence or fragment in each. Perhaps an emoji in one, a GIF in another.

It would be very hard and resource-intensive to apply management controls to these many small messages individually, beyond a single blanket rule. For other reasons that I will get to, I am not sure that searching chat content for keywords would ever reliably find messages that need to be retained, but the question of what you would do if you could make this work is complicated by the fact you are dealing with many, small messages. Would you only preserve the specific message which matched the criteria? Or would you keep some messages before and after? How many messages? From which participants? Can you picture one rule that would work in enough situations that you would consider it trustworthy and reliable?

No way to sort the wheat from the chaff

Chat has no meaningful aggregations or descriptive metadata that would allow us to infer connections between the many messages, beyond the order in which messages were sent and to whom. This would make it very hard to meaningfully apply different rules or controls to different subsets of chat messages.

Each message has some very basic technical metadata automatically attributed to it: who it was sent by, when it was sent, and who received it. There are some tools for labelling a chat or communication channel, but it is incredibly unlikely that such descriptions can be used to help with managing records.

Forever-changing topics, seriousness, and, in the case of group chats, participants

If some content of a chat is of value or well-considered, the chat as a whole is at best like a completely unedited meeting transcript that includes all the banter, interruptions and niceties.

A good meeting record does not look like a complete transcript – and, so, does not look like a chat – for good reason. People do not minute the small talk in the room before everyone has arrived or the jokes made when the meeting is interrupted by a fluffy cat, barking dog, or curious child, for example. Minutes intentionally contain a subset of what was said, focused on the information that is likely to have value for future reference.

This fluid nature of chat is another reason that it would be complex and challenging to apply different rules to different subsets of records. It is as good as certain that any valuable messages will be mixed in with many others that are of no ongoing value, even where metadata is available. For example, we cannot *reliably* say that all chat from one person is more important to retain than another, perhaps because of their role. That one person, if they use chat actively, may write all kinds of things in different contexts, because that is simply how chat works.

Chat is often only one part of a conversation

The information recorded in a chat may have both missing content and context. Quite often it is only one part of the communication between workers.

A chat could easily pick up again after days or weeks of nothing between participants with a message like:

- “Andy found it”
- “Done. Let me know what you reckon. 🙏 or
- “lol” (the evergreen favourite)

The sender may be following on from a verbal conversation, commenting on some event in the participants’ workday like a meeting or an email, or picking up a thread from a different chat. The instant nature of the medium means participants do not need to explain what Andy found, what they have done or what is making them laugh out loud. A conversation may continue at length without even identifying its subject in writing.

Incomplete content and context raise questions about how much chat content could actually be of value. This factor undermines our ability to reuse the information in chats reliably.

What records and information managers could do

A rule of thumb about contractors

When thinking about what to do about a novel format of records like chat, one thing we can ask is what we would expect of the various kind of third

parties that public offices contract to deliver services. The Public Records Act is explicit that a public office is responsible for creating and maintaining records related to matters that have been contracted out. So, as a rule of thumb, any solution we propose to a records problem needs to work equally well for public office employees and for contractors and service providers of any kind.

It is unsurprising that when we come up with solutions we have in mind how we could influence the design of tools and systems within our organisations, where we hopefully have a greater chance of changing how people work than we might expect when contracting services to a third party. However, taking a moment to ask ourselves how a solution would work for contractors can often reveal what matters most and a simpler way of achieving a result that is a good fit for the significance of the problem. Of course, we may choose to take different approaches with employees, individual contractors, and other kinds of service providers, but I think we need a clear, compelling justification when doing so.

So, would we talk to a supplier about what they do with their chat records, specifically? Would we expect them to retain chat? Would we expect them to put in place technology solutions for capturing and maintaining access to valuable chat content over years?

To me, some kind of technical solution for capturing chat records feels like the wrong path to go down. The effort required to manage any chat records of value would be substantial and the benefit in doing so marginal. I would be looking for an alternative means to achieve the same outcome that did not require any substantial technology changes.

I think a public office should always give any third-party direction on what kind of information must be kept and arrange for how to receive such information, as required. In giving such direction, I would be thinking about the purpose of the Public Records Act that relates to being able to hold the government to account. I might consider contractual terms or at least guidance that is relevant to the activity. This already aligns with what Archives New Zealand asks public offices to do.

To be specific about how this relates to chat, in most cases I would not imagine even mentioning chat explicitly. If the above contractual arrangements were agreed, I would accept the residual risk that valuable information ended up in the contracted party's chat messages and nowhere else. That risk feels highly unlikely to come to pass, and it would be even less likely that it would cause a significant problem if it did.

If there was some reason why I thought there was a significant risk that we would not receive the records we need because of something to do with

chat, I might work with the contractor to specify those types of information I think are most at risk. I would set the expectation that the substance of those chat messages needs to be documented elsewhere and might be explicit about where and how.

If this is how I think about approaching the recordkeeping expectations for contractors, I need to have good reason to take a different approach for public office staff. Why would I ask them to do something different or configure the public offices systems to manage the information differently? But let us consider the available alternatives.

We could turn chat off

I cannot picture an office environment where chat will get disabled any time soon. I have read some good arguments for it, but they are as good as moot unless there is another dramatic shift in work culture. The extent to which chat is baked into the incredibly popular Microsoft toolsets would also make this a hard battle to win. And, regardless, you would still need to look at what to do with chat records that are being created now.

We could try to manage different subsets of chat, according to their value and/or risk

For all of the reasons outlined in the previous section, it is not currently practical to identify messages of value, maintain them to a different standard to others, and make them accessible over time. Or at least, we cannot do this in a way that is more systematic and reliable than having conversations with staff about good recordkeeping, about what makes for a good record, and about where and why to save such a record.

Technical solutions may make it easier to identify some key messages in the future, but I am not clear on what a desirable next step would be.

We could keep all chat messages indefinitely

If it is difficult or even impossible to distinguish high-value records in chat, one response is to say we will keep everything and perhaps even protect chat messages from being deleted. This is in line with established practice for, say, a collection of documents, if it is impractical or unhelpful to separate them: retain the whole collection for the longest relevant retention period.

If we did this, we would fail to realise any of the well-understood benefits of disposal and arguably would not meet the requirement on public offices to demonstrate routine disposal.

The organisation would be agreeing to an ever-increasing administrative burden of having to review all chat messages when responding to information requests under e.g., Privacy Act or Official Information Act. Chat's lack of context and often incomplete nature will mean it will get harder over time to successfully interpret information that can be found and evaluate it against a request. This will result in both more effort to complete the request and higher risks of not providing the right information.

Another, perhaps less significant, downside of retaining chat indefinitely is that people may decide chat is a useful knowledge base that they will rely on, ahead of other more suitable options. The consequence of this would only really be a missed opportunity to work more effectively, but that is still worth considering. I noted earlier that people do sometimes choose chat to keep a record of something, but that typically it is of very short-term value. The same reasons that chat results in poor-quality records are the same reasons chat would be a bad choice for a knowledge base longer-term.

There is one real upside to retaining chat that I have yet to mention: chat can be used as evidence for investigating misconduct. Consider bullying, harassment, fraud and what security professionals call "insider risks". Here chat's shift to a medium that unintentionally creates records may be of considerable benefit to an investigator. In my experience we do not normally consider this when appraising information: investigators are often looking for exactly the information that both general staff and information managers think of as "off the record".

There is a tension here. I have mentioned already that the free flow of chat carries high risks that whatever information we retain is hard to interpret or understand later. Flippant, throwaway comments and jokes in a low-context medium add up to a high likelihood that comments could be misunderstood or misrepresented, and this likelihood will go up over time.

So, on one hand, the decision to indefinitely retain chat could mean a public office has access to evidence they would otherwise not have and that they can use that evidence to hold people to account for causing demonstrable harm.

On the other hand, the same decision to retain chat could result in harm for chat participants, for people mentioned, and for the credibility of the public office, if chat is disclosed out of context.

It raises real questions for me about when it would be reasonable to ask someone to defend something they wrote in chat, given what we know about how the medium works. The answer is certainly that *sometimes* it would be reasonable, but certainly not all the time and arguably not even often.

So, I would say that retaining chat forever would only be the best available option in a public office where all of the following are true:

- There is high likelihood of misconduct occurring
- There are very significant consequences of misconduct, and
- Before the introduction of chat there was an unacceptable lack of evidence to identify misconduct and hold people to account

I can imagine that public offices whose functions relate to national security might at least meet those first two criteria, but I think the last point is significant. I would also have to trust that all organisations are working hard on minimising the first point, thus reducing the value of indefinitely retaining chat records.

We could automatically delete chat of a certain age

I have worked in some organisations, private and public sector, where chat messages are deleted after a short period of time, for example 90 days. This is at least possible in Microsoft Teams chats.

This is our best available option, in my opinion. It is getting rid of potentially high-risk and low-value collections of information that were never intended to be treated as records in the first place and requires very little effort to implement. Deleting the records periodically minimises the risks outlined already in this piece around retaining chat.

Arguably, even where there is evidential value in retaining chats in case of misconduct, the likelihood of that evidence being needed for an investigation will decrease over time, so I would still encourage organisations to implement such a policy even if the period for retaining chat is years rather than months.

We have the authority to dispose of “casual or instantaneous communications that contain information of an ephemeral, non-transactional nature”. We can establish a policy internally that affirms that this is the only appropriate use of chat.

This would greatly increase clarity for staff about the appropriate use of tools and about when and where to create records because it takes away the option of ever using chat as a de facto recordkeeping system or knowledge base.

In line with how I might approach contractors, I am not suggesting just turning this policy on, without other controls or initiatives aiming to ensure people can easily do the right thing. I would establish and maintain processes that aid the visibility of the policy to everyone and that aims to ensure people understand

how to work with it. I would include content in staff induction about the topic, ideally tailored to the new starter's line of business, to minimise the risk that people put information only in chat that they need later. In any situation where we are disabling something, we should provide good alternatives, so, in the case of chat, people should know how to capture and share knowledge and significant records and have tools that genuinely help them to do these things.

Once again in line with working with contractors, we may identify particular circumstances in which there is a relatively high risk that valuable information could be lost when chat is auto-deleted, and we can target those sets of circumstances with appropriate solutions. For example, if we found staff who handle interactions with members of the public are less likely to make good records of those interactions outside chat, we could work with those staff to identify the reasons for that and come up with a way to improve their record keeping.

Summary

In this piece I have tried to convey the challenges of chat and my opinions on how best to address them, but I hope this content can be generalised to give you ideas for how to tackle any novel type of records.

To date, I have been in conversations about what technical solutions would be needed to manage chat, have read detailed descriptions of how chat data is held, and so on. I see these conversations as bypassing appraisal and good records management practices. These technical conversations are by necessity about how we manage records not about "what" we are managing or why. They are usually based on an incorrect assumption that we must manage all records to the same standard.

Information and records management professionals are well-positioned to determine the answer to the important question of what information needs to be kept and why. Before chat was commonplace, records managers were not running around public offices telling workers to capture their every thought and water cooler conversation. I believe the content of chat is of the same nature – and value – as those verbal conversations and that our decisions around how to manage chat records must take that into account.

Endnote

1. Archives New Zealand “General Disposal Authority 7” class 1.8.

Bibliography

Archives New Zealand. “General Disposal Authority 7.” 2013. Accessed September 5, 2023, <https://www.archives.govt.nz/manage-information/how-to-manage-your-information/disposal/general-disposal-authorities>