

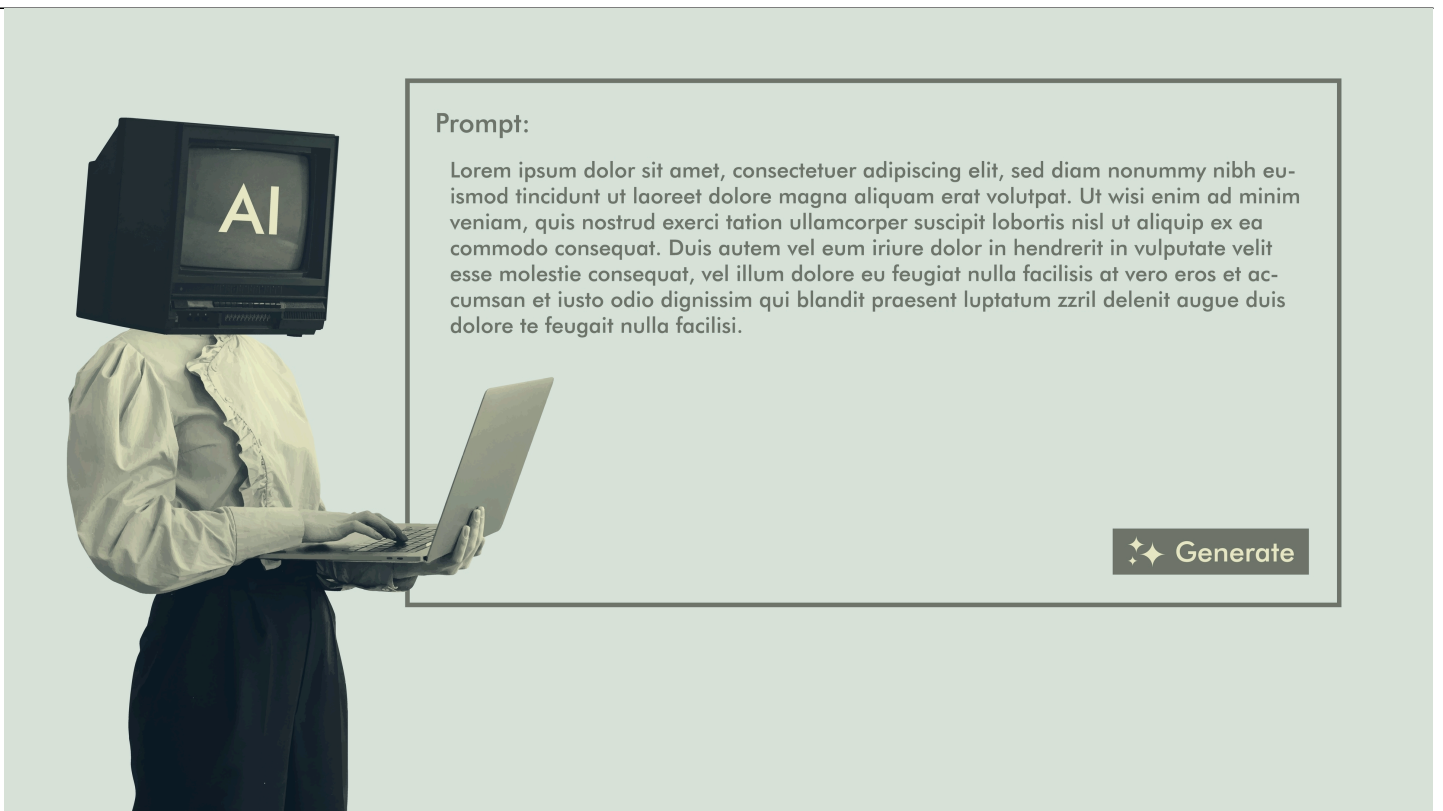


AI Prompts: Keep ‘em or Toss ‘em?

Compliance and Ethics

Law Department Management

Technology, Privacy, and eCommerce



Banner artwork by Collagery / Shutterstock.com

Cheat Sheet:

- **AI prompts are ESI.** Generative AI inputs, outputs, and metadata function like email and other records, making them discoverable and subject to standard legal principles.
- **Retention is a cost-benefit decision.** Organizations must weigh litigation value, institutional knowledge, and reuse against risks like discovery exposure, storage costs, and data breaches.
- **Shorter retention is often appropriate.** High-volume, ephemeral AI prompts typically warrant retention periods measured in days or months, especially when sensitive data is involved.
- **Protect high-value prompts separately.** Prompts with competitive or proprietary value should be stored outside AI systems and managed like trade secrets, not left in default platform histories.

Imagine a technology that everyone in your company puts information into every day. Their input includes text and other media. It runs from the mundane to the highly sensitive. And once turned over to the technology, the input resides in servers.

Keeping such a wide variety of data on servers has plusses and minuses. Often the stored info helps your company defend lawsuits. Other times the retained data has created liabilities. To manage these costs and benefits, your organization sets rules on retaining and deleting data generated by the technology.

The tech I'm talking about is, obviously, email.

AI prompts are just electronic data



Source: Author, ChatGPT

Of course, the description above doesn't apply to *just* email. Generative AI technology — like ChatGPT, Claude, or Google Gemini — also takes in user input and stores it on cloud servers.

With Gen AI the input is usually called “prompts.” Here are a couple: “Make an image of a chow chow paddling a canoe” and “Summarize *Heated Rivalry*.”

Along with prompts Gen AI platforms capture outputs (for instance, AI-generated text or images) and metadata (like activity logs).

From a data management perspective, email and Gen AI share a fundamental similarity: both technologies generate and store lots of ones and zeros. And for in-house lawyers crafting document retention policies or managing e-discovery, that means massive amounts of [electronically stored information \(ESI\)](#). AI prompts are at bottom just a subset of ESI.

Courts recognize as much. Take the recent ruling in *U.S. v. Heppner*, No. 25-CR-503 (S.D.N.Y. Feb. 17, 2026). Judge Rakoff of the Southern District of New York [held](#) that a litigant’s communications with the large language model (LLM) Claude weren’t protected by attorney-client privilege or the work-product doctrine. In other words, the prompts were discoverable ESI. “AI’s novelty,” the Court concluded, “does not mean that its use is not subject to longstanding legal principles.”

In the ongoing multidistrict litigation against OpenAI for copyright infringement, Magistrate Wang of the Southern District of New York [reached a similar conclusion](#). In December 2025, she ordered OpenAI to produce 20 million output logs from ChatGPT, applying the same Rule 26 analysis used with any other ESI. *In re OpenAI, Inc., Copyright Infringement Litigation*, No. 25-MD-3143 (S.D.N.Y. Dec. 2, 2025).

Yet, despite this essential sameness of prompts and emails, many organizations approach prompt retention as a unique “AI” issue rather than as a document retention matter.

But like the courts, organizations and their in-house counsel should think about how to retain and delete prompts the same way as they would any other ESI.

That means, most importantly, weighing costs against benefits.

But like the courts, organizations and their in-house counsel should think about how to retain and delete prompts the same way as they would any other ESI.

Trade-offs

Some documents’ retention periods are driven by law. Under the [FLSA](#), for instance, employers must keep payroll records for at least three years. Tax records are commonly kept for seven, the [limitation set by the IRS](#) for claiming credits or refunds for bad debts and worthless securities. (So, hang on to your [\\$HAWK token](#) receipts for a few more years.) Common law duties and the [Federal Rules of Civil Procedure](#) also require preserving documents, including ESI, relevant to litigation.

But the law doesn’t dictate a retention period for every type of record. In fact, for the larger volume of data generated within a company — emails, instant messages, Word docs, and more — companies

need to decide on an appropriate retention period.

That decision boils down to a basic question: At what point are the benefits of retaining a type of document outweighed by the costs? Once it goes past that point, toss the doc into the abyss.

Benefits of retention abound. Retained docs can help the organization pursue or defend litigations and claims. They spare employees from recreating work product from scratch. Certainly in-house lawyers — shh! — routinely repurpose advice given in prior memos and emails. Keeping records is especially important in organizations with frequent employee turnover; documents may be the only place where the company's institutional knowledge lives.

The costs of retention may be less obvious, at least to non-lawyers. Just as documents may help pursue or defend litigations, they can also create litigation exposure. Every litigator has seen their share of “bad” emails — well, “emails that might be misinterpreted.” But even if the records are harmless, keeping them still imposes costs. That many more retained records means that many more records open to discovery or subpoena. A company that deletes emails after six years may eat millions more in discovery costs than an organization that deletes emails after six months.

But put legal concerns aside, records and ESI also take up space. That not only increases hosting costs but also poses a security risk: the more data a company keeps, the more data may be compromised.

[Check out the ACC AI Center of Excellence for In-house Counsel today.](#)

AI prompts — keep them short(ish)

It's easy for an in-house lawyer to list the costs and benefits of keeping docs; weighing them against each other is harder.

There's no objectively “right” approach. Companies must take into account the nature of their industry, their business needs, the data at issue, and more. A company in a lightly regulated field that rarely touches sensitive information may see little harm in fixing a three-year retention period for emails. A business handling medical data may clear its electronic correspondence faster than a budding politician deletes their old tweets.

Companies must take into account the nature of their industry, their business needs, the data at issue, and more.

So even as approaches vary, some rules of thumbs still apply. Ephemeral, high-volume records — like emails or Slack messages — generally warrant a shorter retention span. Low volume, frequently referenced materials — say, contracts and tax documents — should be kept longer. The more sensitive the data, the shorter the time span too.

So where do AI prompts sit on that spectrum? Well, they're probably closer to Slack messages than

tax records. A single AI-generated image may result from hundreds of prompts. Those interim prompts have some value; you might retrace your steps, to recreate similar materials or redo your work.

But your prompt history grows less useful over time, while presenting the same discovery and data security risks. A half-year out, is it really doing much? And if the prompts include personal information or other sensitive or regulated data, all the more reason to consign it sooner to the digital dumpster.

Not surprisingly, [Google](#) and [OpenAI](#) allow enterprise admins to set short, even zero-day, retention periods. If you're looking for a default — and assuming no legal or contractual requirements to keep materials — think days or months, not years.

Some companies may nonetheless choose to hang onto prompts longer, perhaps to defend from potential claims. For instance, an org might point to its prompt history to argue that its AI-generated photo didn't copy another work: "See, we never asked for [Guernica](#)." But before taking that stance, you should be confident that your prompt history exculpates more than inculpates. It can always hold a smoking gun: "Computer, make me a *Guernica* lookalike. But more cheerful."

Even if your goal for retaining prompts is to defend against copyright claims, it's hard to imagine a retention period longer than three years — that is, [the statute of limitations for a copyright infringement claim](#) (though [there is nuance to that deadline](#)).

Proprietary prompts

Wherever you land on a retention policy, document deletion shouldn't mean ditching prompts with competitive value. Some organizations treat prompts as a proprietary secret sauce. "We can make canoe-paddling chow chows that our competitors only dream of!"

But if that's you, don't rely on the AI tool to house those prompts. Instead store them in a separate location; and if appropriate apply [the same procedures you would to trade secrets](#).

Conclusion

Ultimately, the retention period you set doesn't need to be perfect. Generative AI tools remain relatively new, and you can always change your policy as the technology and industry norms evolve. What's key is to approach prompt retention thoughtfully and consistent with your organization's overall approach to ESI.

Now delete this article after reading.

[Join ACC for more AI guidance.](#)

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a

definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

[Christopher Wlach](#)



Senior Director, Legal

Acxiom

[Chris Wlach](#) is the senior director, legal of Acxiom. Before moving in-house he focused on complex commercial litigation at Arnold & Porter. He is a certified information privacy professional (CIPP/US) through the International Association of Privacy Professionals. He also chairs the board of [HEART](#), a humane education nonprofit.

