



2024 Law Firm Information Governance Symposium

AI Considerations for IG Processes



Contents

04/ Introduction

05/ Privacy and regulatory compliance

05/ Client directives

07/ Intellectual property concerns

08/ Matter mobility

09/ Retention and disposition

10/ Mandated destruction

11/ Technology and data governance; securing information

12/ Third-party engagement

12/ IG awareness and education

13/ Preservation and auditing

14/ Matter lifecycle management

15/ Conclusion

16/ Appendix A

16/ Appendix B

Authors

Elizabeth Adkins

Director of Information Governance, Davis Wright Tremaine LLP

Karen Allan

Information Governance Manager, Wiley Rein LLP

Brianne Aul

Senior Manager of Information Governance Strategy and Operations, Morgan, Lewis & Bockius LLP

Bryn Bowen

Practice Group Leader, Compliance, Intapp

Colin Cahill

Senior Manager, Information Governance and Collaboration, InOutsource Consulting LLC

Scott Christensen

Founding Member, Kyndig Consulting

Galina Datskovsky Ph.D., CRM, FAI

Director, OpenAxes

Brian Donato

Chief Information Office, Vorys, Sater, Seymour and Pease LLP

Michele Gossmeier

Director of Information Governance, Holland & Knight LLP

Christine Hunter

Senior Counsel, General Counsel's Office, Director, Information Governance, Blake, Cassels & Graydon LLP

Jessica Marlette

Information Governance/Data Strategy Attorney, DLA Piper LLP (US)

Rishi Maharaj

Associated Director of Records and Information Governance, Cravath Swaine and Moore LLP

James Merrifield

Director of Information Governance & Business Intake, Robinson & Cole LLP

Reggie Pool

CEO, Practical IG

Charlene Wacenske

Director of Record & Information Governance, Morrison & Foerster LLP

Leigh Zidwick, CIGO, CIP

Senior Director, Information Governance, DLA Piper LLP (US)

Introduction

The term artificial intelligence (AI) was originally coined in 1956 and has evolved over time from expert systems to machine learning (ML) and natural language processing (NLP) and now to large language models (LLMs) and generative AI (GenAI). While the use of AI in various forms is not a novel concept for law firms, the adoption of GenAI could revolutionize the practice of law, offering unprecedented levels of efficiency and innovation (see Appendix A for a more detailed explanation of terms). Unlike its predecessors, whose functions are primarily analytical, GenAI can create original legal content. As described in our previous paper on [Generative AI and Law Firm Information Governance](#), these expansive new functions create new risks. It is imperative for firms to approach GenAI tools with a measured understanding of their capabilities, limitations, and risks, including how the technology impacts an organization's information governance (IG) program and vice versa.

This paper will build upon the IG-specific considerations highlighted at a high level in our [previous paper](#) as of July 2023. Here, we delve into more detail as to how GenAI impacts and is impacted by certain law firm IG process components as outlined in the "wheel" in our [LFIGS 2.0 paper: An Established Law Firm Information Governance Framework](#). A complementary paper will follow on the technological aspects of Generative AI through the lens of IG. The goal is to update it periodically as significant technological changes occur in the AI/IG space.



Privacy and regulatory compliance

One of the key challenges in using GenAI for legal work is ensuring that such use complies with applicable laws and regulations, including privacy law. [A January 2024 study by KPMG](#) found that 63% of consumers were concerned about the potential for GenAI to compromise an individual's privacy by exposing personal data through breaches or other forms of unauthorized access or misuse. The fact that the legislative framework is nascent makes IG professionals' roles even more challenging.

Countries are designing and implementing GenAI governance policies and legislation at varied paces. Efforts include the development of comprehensive legislation, focused legislation for specific use cases, and voluntary guidelines and standards. There is no standard approach, and the democratization of access to LLM tools is heightening the importance of having something over nothing. Given the transformative nature of GenAI technology, the challenge is to find a balance between promoting or allowing innovation and regulating risks. Therefore, governance of AI often begins with a jurisdictional strategy or ethics policy instead of legislation.

Many current legislative efforts are focused on regulating GenAI technology providers. For example, in May 2024 the EU passed its [Artificial Intelligence Act](#) (with overwhelming support) to establish a uniform framework for the development and marketing of AI systems that are considered 'trustworthy' and ensure the rights enshrined in GDPR. In October 2023, the United States issued guidance to help establish guardrails surrounding AI development.

Law firms are only indirectly affected by the above; what they need to be most concerned about is professional regulations that set out conditions for the use of AI in providing legal services to clients. One key topic that existing professional guidelines cover is transparency around the use of GenAI. The California and Florida Bars were the first state bars to issue draft ethical guidance on GenAI, and several other state bars and the American Bar Association are working on them, with more and evolving guidance to come. These guidelines should inform each firm's ethical principles, which in turn can be shared in response to client requests. It's advisable for lawyers, and by extension, IG professionals to stay informed about jurisdiction-specific regulations and professional guidelines to determine whether disclosure or other steps are necessary prior to using GenAI in legal work.

Client directives

In addition to the regulatory framework and guidance discussed above, law firms must also consider their contractual arrangements with clients. As outlined in our previous paper, [Practical Solutions to Implement Client Information Governance Requirements](#), client guidelines are a key part of client engagements, and as such, utilization of GenAI should also be part of these discussions. Clients are increasingly including clauses in their outside counsel guidelines (OCGs) that impose restrictions or conditions around the use of such tools. Some clients may understand the risks versus benefits of GenAI better than others. Some may end up opting out of including their data in any GenAI tool or model or extending their concerns not just to GenAI, but to all forms of AI. The scope of the OCG requirements can vary widely, but examples of OCG terms received by some firms to date include:

- A requirement for consent or written notification prior to using the client's information in connection with GenAI tools or applications.
- A prohibition on using the client's information to train the tools unless safeguards are in place to protect the information from unauthorized access. (Some clients ban the use of the client's information to train the tools altogether.)
- Requirements to supervise and review any work product created by leveraging GenAI to address concerns about hallucinations.
- A requirement to ensure that the firm has the ability to permanently delete client data input into the GenAI tool(s).

- A prohibition on the use of public tools (e.g., the public version of ChatGPT).
- A requirement to add a header to any deliverable produced using GenAI tools and disclose the usage in the invoice for that work.
- A list of prohibited uses, such as using GenAI tools for drafting patent applications or for performing prior art searches. Alternatively, some clients may list acceptable uses, such as using the tool to summarize case law or to generate boilerplate contractual clauses.

- Requiring the firm to represent and warrant that content generated using GenAI models does not infringe on third-party intellectual property rights.

The number and variety of requirements may prove challenging to navigate and add another layer of complexity to creating a holistic and strategic approach to the deployment of GenAI across the firm. IG professionals, together with firm leadership, will need to determine whether and how to negotiate these clauses, and should be prepared to share with clients the firm's vision for how they will work with GenAI.

It is advisable to create a GenAI committee within the firm to navigate and respond to clients' concerns with respect to their information.

As tools become more widely acceptable, the angst about the use of GenAI may subside. Ensuring transparency in every aspect of the selection, use, output, and monitoring of GenAI tools will be a critical factor in addressing clients' concerns and encouraging them to accept GenAI use as an integral part of providing legal services. It is advisable to create a GenAI committee or working group within the firm to navigate and respond to these concerns with respect to their information. Clients should know that their concerns are being carefully considered by the firm. At a minimum, the committee or working group should include representation from IT, the General Counsel's Office, Information Security, Privacy, Knowledge Management, and Information Governance, as well as lawyers (including partners and associates) who are advising clients on their use of GenAI, or who are advocates for the use of AI in improving accuracy and efficiency of client service.

It is also advisable to establish a set of principles on the appropriate use of GenAI within the firm, based on applicable legislation, professional guidelines, and ethical and legal responsibilities to the client. These principles will inform the policies around GenAI use, which can be incorporated into a firm's existing Acceptable Use

Policy or inserted as a standalone policy. The GenAI policy should address the process for requesting the approval of any new GenAI tool for use at the firm, as well as requirements for the use of approved GenAI tools (referencing a published list of approved tools that can be easily found by all personnel.) Examples of risks to the firm, and how those risks inform the unacceptable uses of GenAI, should also be included in this policy.

Some clients are very interested in innovative technologies and may want to partner with the firm to pilot or develop GenAI technologies using their data. They may expect the firm to invest in GenAI and want to know about that investment. In this case, transparency also serves to demonstrate the firm's commitment to legal technology innovation and client service.

It is worth noting that several vendors are incorporating AI capabilities (predominantly based on ML) into their products that are designed to track and extract terms from OCGs, making the process of reviewing, tracking, and complying with client requirements much more efficient.

Intellectual property concerns

There are three key intellectual property questions to consider when using GenAI that may impact a law firm's IG function:

1. Who owns the data that is input into the tool?
2. Who owns the data generated by the tool?
3. What risks to the client's intellectual property are introduced when using GenAI tools?

The answers to these questions are the subject of emerging debate, litigation, and regulation.

Under US law, patents and trade secret rights are at risk if exposed to the public. Even a seemingly minor exposure can have major consequences on the ability to maintain protection and ownership of these valuable assets. For these reasons, when working with patent or trade secret records, it is particularly important to use only private instances of GenAI tools, meaning instances that are accessible only to a particular firm, with data that is confined to that firm's instance of the tool, protected from unauthorized access, and not commingled with the data of any other entity. It is also essential to make sure that the client's intellectual property is not used to train a GenAI tool beyond the confines of that client's data. As mentioned above in the section on client directives, even with a closely protected and controlled private deployment, some clients may wish to prohibit the inclusion of their data as a risk-avoidance strategy.

Most GenAI platforms rely on a model that has been trained using publicly accessible data, which allows them to be deployed within a firm's environment with predetermined exemplars and best practices as a starting point for possible internal training. The major foundational LLM developers have ingested huge volumes of data from the internet and other public sources, which has raised concerns from authors, artists, and programmers about lack of consent or compensation for the use of their works in GenAI tools. As of this writing, more than a dozen lawsuits have been filed seeking damages for what the plaintiffs see as a violation of copyright laws.

Technology companies argue that copyright laws allow usage under the concept of "fair use." They point to

a copyright infringement suit brought in 2005 by the Author's Guild in response to Google's initiative to manually scan millions of books to create a searchable database that provided portions of the books in response to searches. In this case, a federal appeals court ruled that Google's use of published works was fair use due to the transformative nature of the database being created (i.e., the benefits outweighed the risks). Although the companies have won some early rulings, it's far from clear whether the courts will find that training AI platforms based on copyrighted works without permission is a violation of law.¹

The uncertainty of the legal permissibility of using copyrighted sources to train GenAI has led the major GenAI providers to issue indemnities to help protect and reassure their clients. GitHub led the way in June 2022 by agreeing to indemnify users of its Copilot coding tool from infringement claims. In September 2023, Microsoft announced the "[Copilot Copyright Commitment](#)," which was later renamed the "Customer Copyright Commitment." As part of this commitment, Microsoft will defend any customer challenged for violating copyright by using its platform and will pay any penalties that result from legal action. OpenAI announced a similar initiative with its "[Copyright Shield](#)" in November 2023.² It's important to note that all of these indemnities are only effective if the customer has followed the provider's guidelines on how to responsibly use their products.

In December 2023 *The New York Times* filed a copyright infringement lawsuit against OpenAI and Microsoft, contending that not only were millions of its published articles used to train automated chatbots without a license, but that the response to chatbot prompts yields nearly verbatim *New York Times* articles. The lawsuit refers to these prompt results as "regurgitation." OpenAI has rebutted that regurgitation is an unintentional "rare bug," and that the *Times* "intentionally manipulated prompts" to get it to happen.

The New York Times lawsuit and similar lawsuits may be settled in part by the payment of licensing fees to the copyright owners, but in the meantime, AI providers are facing some potentially large judgments against them.³

More open questions remain about IP rights to GenAI output. The US Copyright Office, the US Patent and Trademark Office, and European regulators have declared that only a human can be an author or inventor (although other countries, including China, have stated otherwise). There is some question as to what this means when a GenAI tool generates new content based on human prompts and queries. Is it possible for a GenAI platform to be a co-inventor or co-author? Currently, the US Copyright Office does not permit registration of any works that were generated with AI, but it's not clear how the US Patent and Trademark Office will rule on this.⁴

While the question of ownership of output might be more of a concern to those who are generating creative content or programming, there may also be considerations for law firms using GenAI to create client work product. For example, could the client conceivably own not only the work product but also the prompts and queries that created the work product? Will clients look to their law firms for intellectual property indemnities to protect them in the event their content is incorporated into a model and inadvertently used by a third party, or if the work product they think they own free and clear is threatened by a third-party claim? Firms should consider the extent to which their inputs will be used to train

the model, and ensure they adequately address these risks in their contracts with GenAI service providers. Firms should think particularly carefully about the idea of generating text in the voice of a particular individual because that could also yield thorny issues around ownership. For example “write a draft complaint for X in the voice of Alan Dershowitz.”

With so much ambiguity around intellectual property rights and GenAI, many see a need for regulation to make the path forward clearer. [The EU Artificial Intelligence Act](#) described above provides some certainty: it requires providers of general-purpose AI models to publicize summaries of the content used for training.

In the United States, federal guidance for GenAI relevant to IP was issued by the US Patent and Trademark Office and the US Copyright Office.⁵ Several states to date are in the process of drafting their own AI regulations, including California—but much of the initial focus addresses bias, cybersecurity, and privacy, not intellectual property.⁶

The bottom line: IP considerations in GenAI are in flux, not likely to be resolved soon, and will likely vary from state to state; they should be monitored closely.

Matter mobility

Matter mobility refers to transfer of a client/matter from one firm to another, whether with a transferring lawyer or otherwise, and is discussed in our earlier papers, including [Matter Mobility 2.0](#). In addition to the usual requirements that must be considered to comply with matter mobility policies and best practices (including ethical rules, client consent, contractual obligations, potential conflicts of interest, and the type of practice), IG professionals now need to consider how GenAI may impact these processes.

Recent case law has emerged related to the proprietary knowledge and skills that certain parties may develop about GenAI and how this in turn may impact individuals' moves from one organization to another. For example, a US court recently considered a transferring lawyer's knowledge and skills in a matter involving GenAI technology in hearing a motion by his former firm to disqualify him from continuing to act for that client

at the new firm. In another case outside the law firm context, the court considered transferring an executive's AI knowledge and skills as a factor in evaluating the enforceability of a non-compete clause (see Appendix B for a summary of these cases). Another mobility-related point to consider is that depending on the jurisdiction, a law firm may or may not include drafts when it transfers a file. If your firm transfers drafts, consider whether any additional steps are needed to label AI-generated drafts of legal documents prior to transferring them. To meet professional standards, lawyers must review and validate any AI-generated work product. If a transfer happens before that review, the best practice would be to clearly identify any AI-generated work product before transferring so that it can be reviewed by the new firm. More generally, prior to transferring, lawyers should consider reviewing whether any additional disclosure or client consent requirements need to be met.

Retention and disposition

IG professionals at law firms with more mature retention and disposition programs have likely already adjusted their policy, schedules, and processes to accommodate new repositories, including those repositories that originated as or migrated to a software as a service (SaaS) platform. However, records created via GenAI add new complexity for retention programs and disposition

processes because the impact of disposition may not **just** be a given record itself but also any new records generated based on it. Depending on the technology used, it could be very difficult, if not impossible, to gather all such records, especially as the scope of newly created content proliferates.

Organizations need to consider the fundamental question “What is part of the record?” when determining what to retain/dispose related to AI.

Thankfully, even as GenAI proliferates, many of the more traditional questions regarding retention processes remain the same and can help frame the discussion. For example, an IG professional should consider the following:

- **How long is the firm legally required to retain the source records used to train, fine-tune, or augment the model?** The IG professional should raise this question with their General Counsel or Compliance team as soon as possible for incorporation into the firm’s larger GenAI strategy. While legal and regulatory requirements regarding the retention of GenAI content may not yet be clearly defined, keeping tabs on this question will ensure that it is part of larger retention strategy considerations as new solutions are vetted and implemented in accordance with the firm’s GenAI roadmap.
- **Is there a business requirement to retain the source records longer?** This might be one of the most critical discussion points for the IG professional to revisit. One of the key benefits of GenAI is its ability to learn and generate from incredibly large data sets, in a way that is impractical or impossible for a human to do. As such, a prior decision from a department or practice group that a “lookback” period is “X” period of time may be reconsidered through the lens of machine learning—what may be a reference or scope limitation for a human may no longer be a hindrance with AI technology. For example, should a document’s normal retention period be extended now that it is used for purposes of “feeding” the AI technology? If so, the IG professional should ensure those documents or repositories are properly tagged in accordance with their policy.

- **Alternatively, is over-retention actually degrading the return on investment (ROI) on current and future GenAI initiatives?** While there could be justification for a longer retention period on certain data, just as easily one could argue longer, and at times unnecessary, business-related retention schedule rules impede quality GenAI outputs, especially if the input data contains information that is otherwise irrelevant, outdated, and/or is generating responses that are not useful to current and future practices. Policies are a good example of this, where older versions may feed into confusing or inaccurate results. Beyond regulatory and client requirements, firms with longer, indefinite, or non-existent retention periods now face additional ROT (redundant, obsolete, trivial) data challenges and may find that the ROI on GenAI is diminished by the potential “garbage in, garbage out” risks that over-retained information poses. Firms with knowledge management (KM) or similar departments can be especially helpful here, as there may be opportunities to identify certain sets of documents already tagged as precedent or know-how, and likely key resources for future AI solutions. Absent a KM department or solution, IG professionals can also revisit retention decisions with practice groups and administrative departments to whittle down retention periods or the data sets they apply to in a much more manageable (and AI-beneficial) manner. AI may be the perfect driver to focus on retention policies that are overdue for review.
- **What client requirements are in place regarding retention?** Beyond OCGs, which may include rules around how clients’ data will be used in any GenAI

tool (see above), the IG professional needs to consider the impact of GenAI tools on client-mandated retention (particularly when client requirements conflict). For example, if a client expects all data to be returned or destroyed within X period of years, and simultaneously expects the firm to use GenAI to provide more efficient legal services to them, does that mean certain documents are now expected to be incorporated into an AI tool and therefore exempt from their retention requirements? Identification of those clients with possible discrepancies in their instructions will help ensure the firm is not agreeing to conflicting mandates and/or revisiting the conversations with those clients to determine if revisions need to be made to the guidelines themselves.

- **How can data be disposed at the end of the retention period?** As mentioned, SaaS-driven solutions have already forced IG professionals to reconsider who is responsible for retention—whether policies and processes can be set within a particular tool vs. what may need to be purged by a vendor. Expansive usage of GenAI solutions will make it significantly harder—if not impossible—to remove input data, whereas private solutions are easier to manage (though not without impact to future generative work product). The GenAI model may have an ephemeral retention policy, i.e. the data input into the model is discarded from memory following output generation; however, the solution

elements elsewhere in the data flow that carry the input to the model and transform the model's response into solution output for the users will have their own retention policies, which may vary at each step of the data flow. In most instances, only the vendor will have access to these data repositories, and destruction of data within them on a set schedule may be a contractual requirement of the vendor and not a task for the IG professional. These retention concerns and the evolution of GenAI tools and their related contracts merit regular review of firms' policies and guidelines regarding use (or restrictions on use) of public GenAI solutions. We will examine this and other technology concepts in future papers as referenced in the intro.

Again, retention and disposition questions in the context of GenAI systems are not novel; the responses, however, may be unique and may vary depending on the system.

It is important to note that AI can also be used to support a firm's IG function by tagging data in a way that helps bolster its disposition program. While not necessarily GenAI, existing AI solutions can identify materials that could otherwise pose a retention and/or security risk to the firm, such as data that contains sensitive information or personal data like an individual's birth date, SSN, credit card number, or health information. As IG continues to grapple with GenAI, it's important to remember that there are multiple opportunities among the challenges.

Mandated destruction

Most destruction processes are the same for standard disposition and mandated destruction. However, GenAI adds new considerations that must be addressed by IG professionals, together with the firm's General Counsel. For example, if GenAI was used in a litigation matter for which a destruction request is issued, it is important to understand the extent to which the GenAI records are subject to destruction. For example, are the prompts and inputs subject to destruction along with any outputs (our previous point around 'What is the record?')? Further, it will be critical for the firm to understand if and how data within a cloud-based GenAI tool can be identified and destroyed. Depending on the tool, the firm may be able to destroy the data via its admin privileges or may need the vendor itself to assist with the purging of content. Firms should obtain contractual commitments as to a vendor's capability to delete data and the turnaround times for

such deletion requests prior to onboarding them or as part of the contract renewal process with the vendor.

Mandated destruction requests may permit an archival copy that may be subject to the same access restrictions as the relevant records to be retained, generally stored in a firm's document management system for the duration of the firm's retention period. GenAI adds an extra layer of complexity to this concept. For example, can a firm retain a copy of records that exist in a GenAI system (or, even more complex, a copy that has been used to train or fine-tune a GenAI model)? If the record was uploaded as part of a prompt to a GenAI model, it may still need to be removed to prevent future use. However, if the record has been used to train the model, it will not be possible to "remove" or reverse the resulting enhancements to the model. IG professionals will need to ensure they account for this as part of their process.

Technology and data governance; securing information

Law firms must apply strong information rights management (IRM) practices to help manage conflicts of interest, preserve client confidentiality, and ensure their personnel can easily access the information they need to do their jobs well. This is when implementing any GenAI technology because such systems possess the capability to sift through extensive datasets. Without stringent IRM, there is a risk that GenAI systems could inadvertently expose confidential information to unauthorized individuals.

A law firm that develops or implements a private GenAI tool with broad access to data within its systems needs to be aware of the risk of data leakage and inadvertent disclosure. When designing any such tool, IG professionals need to ensure that access rights to the training data are properly applied, that training data is sufficiently anonymized where needed (and, if applicable, that the tool can effectively differentiate between anonymized and non-anonymized data).

One specific component of IRM that is ripe for both GenAI risks and opportunities—and relevant to IG professionals—is the implementation of ethical walls. The importance of appropriate ethical walls is amplified in larger law firms in which numerous teams handle multiple matters concurrently. Some firms see the potential to design customized GenAI tools for the purpose of identifying conflicts of interest and implementing ethical wall requirements. Here is an outline of how this might work:

- **Conflict identification:** A GenAI tool could be developed and trained on a firm's conflict of interest policies and past conflicts data. When a new matter arises, the tool could analyze relevant details (e.g., party names and nature of the matter) and flag potential conflicts. This would help lawyers quickly identify potential conflicts early in the intake process. It's important to note that the tool would need to be able to "show its work" by referencing the data on which it based its response.
- **Data screening and anonymization:** GenAI could be used to screen documents and emails for sensitive

information before they are shared with lawyers working on different sides of an ethical wall. The tool could identify and anonymize relevant data points, mitigating the risk of inadvertent conflicts.

- **Ethical wall training:** Once trained, the tool could be used to develop and deliver training on conflict-of-interest rules and the firm's ethical wall policies. Lawyers could interact with the tool through simulations or Q&A sessions to better understand how to identify and avoid conflicts.

It is important to remember that no GenAI tool can replace human judgment. This is especially important in the context of identifying and enforcing ethical walls due to the high risk involved in law firm conflicts of interest. Here are some additional IG-related points to consider:

- The effectiveness of any GenAI tool depends on the quality of the data it is trained on. The training data should be comprehensive and up to date on the firm's ethical walls and relevant case law. Consider risks of bias in training data.
- Robust anonymization techniques should be implemented and reviewed regularly to ensure that any tool interacting with conflicts data does not surface data that users should not see.
- Lawyers should always review all outputs and exercise their professional judgment when making decisions about conflicts of interest.
- As with any technology, strong data security protocols should be implemented to safeguard information.

By acknowledging these concerns and taking appropriate steps to mitigate them, law firms can leverage the private LLM as a valuable tool for enforcing ethical walls while still upholding their professional obligations. Remember, any LLM should be seen as an assistant, not a replacement for a lawyer's ethical judgment.

Third-party engagement

New technology services should be properly assessed for risk by the law firm before being implemented. AI tools are no exception and may require an additional level of scrutiny, due to the heightened level of risk posed by this quickly advancing technology. In addition to the standard technology risk assessments that your firm performs, other areas of consideration for AI vetting include:

- Ownership of your data: Prevent giving up rights, title, or interest to your data. Pay particular attention to whether your data can be copied, stored, displayed, or processed without your consent.
- Understanding how both inputs and outputs are handled and the extent to which they are used to improve the AI services, train models, or other purposes, including the potential for unintended use.
- Understanding whether the AI vendor will share your data with their third parties and, if so, how the data will be protected, retained, and ultimately disposed of.
- How the vendor will protect and manage your confidential information from ingestion to final disposition, and how well the agreement with the vendor guarantees these custodial duties will be fulfilled.

IG awareness and education

Lawyers and IG professionals need a strong understanding of GenAI tools and how they impact information governance. Lawyers should carefully review and verify the accuracy and completeness of drafts, checking for errors, omissions, and inconsistencies, ensuring they comply with applicable laws, regulations, and ethical rules and that they do not infringe on any third-party rights. If not already covered in the client's OCS, lawyers should understand client expectations or obtain the informed consent of the client before using AI technology in the generation of drafts or other forms of work product. Part of that discussion should address the specific needs and circumstances of the case, and the lawyer needs to explain the benefits, risks, and limitations of the technology.

IG implications of GenAI

The privacy, security, confidentiality, and IP aspects of GenAI have been the focus of the legal industry's education and awareness efforts since late 2022. For instance, see page 7 of our 2023 paper, [Generative AI and Law Firm Information Governance](#). It's not surprising that navigating these information-related risks has been a key focus for law firms, law schools, and industry groups. Whether via a formal policy or general communication/education, when ChatGPT began proliferating in late 2022 and early 2023, many law firms promptly cautioned their lawyers not to rely on any GenAI-based systems

for substantive legal work or to enter any confidential information into any such platform while they embarked on a mission to understand these tools and their potential benefits and risks.

Similarly, law school administrators had to grapple with issues of academic integrity early on: how can professors identify when a student has relied on ChatGPT or a similar tool to the extent that it constitutes plagiarism; and in what circumstances should students be penalized for such use? Many universities have updated their academic dishonesty policies to prohibit the use of GenAI-produced text unless explicitly permitted by the instructor. To add complexity to the task of enforcing such policies, tools that have been created to identify non-original text can also be used by students to try to evade AI detection.

Building on these initial directives, law firms and law schools alike are keen to ensure that their lawyers and students fully understand how information will be treated when they input it into a GenAI system. Cautionary tales were circulated, including the notable case of [the New York lawyer who got caught](#) relying on fake cases that ChatGPT had hallucinated in front of a federal court judge. More recently, a business owner representing himself in an unpaid wages lawsuit was sanctioned by the [Missouri Court of Appeals for writing legal briefs](#) containing almost two dozen fake case citations.

Law school curricula are being revised to incorporate similar learning objectives around confidentiality, security, and accuracy. Such considerations can be built into existing courses on legal ethics and professionalism, IP, privacy, legal research and writing, information technology, access to justice, and certain international law courses (including the cross-border nature of information governance issues).

Law firms should continue to prioritize educating their lawyers and staff on best practices to protect their firms' and clients' information. A law firm's General Counsel's Office and IG professionals play a key role in communicating the firm's GenAI policies, client directives related to GenAI (see the discussion on OCGs above), and the risks of non-compliance with such rules. This shift to incorporating GenAI into the firm's training syllabus is particularly important because many lawyers are already using public GenAI tools in their daily lives or may be experimenting with GenAI in their professional lives.

Law students have been using GenAI tools to generate first drafts of case summaries, cover letters, resumes, and research memos, as well as to revise emails. As they enter law firms, they will expect to continue using such tools to support their education and, ultimately, their legal practice. Overall, this proliferation emphasizes the need for internal guidance and governance around the use of GenAI.

Available GenAI tools and technical skills

Understanding the menu of available technologies, including whether each relies on "traditional" AI or GenAI and how each transforms inputs into outputs, is imperative to identifying what use cases may be appropriate in any given context and identifying potential IG and other risks. For example, the GenAI-supported case summarization capabilities of legal research services do not raise the same information governance

issues as the use of GenAI to summarize client documents uploaded into the system.

Technical skills such as prompt engineering/problem formulation are currently very helpful, but the relevance of these skills may change as tools become more advanced and as GenAI components are built into existing tools. We will examine this and other technology concepts in future papers.

Many law firms are designing curricula for their lawyers and staff on GenAI, whether through foundational workshops for the entire firm population or training customized to a particular practice group or business services group's potential opportunities and risks. Law firms can also learn from law schools, which, as noted above, are building GenAI content into existing courses as well as designing new courses and extracurricular offerings focused on legal technology, of which GenAI will form a component. Some forward-thinking professors are building GenAI elements into seemingly unrelated courses, reflecting the reality that this technology will become ubiquitous. For example, a University of Toronto law professor and co-founder of a legal tech company has incorporated the company's predictive tax law tool into the school's tax law courses. Osgoode Hall Law School is exploring GenAI uses in a number of courses, including a new micro-course in AI and Technology in Legal Practice, to cover the technology skills students will need as they transition into practice. Several of these courses are taught by practitioners, evidencing a translation of what students will need in practice into early law school education.

As noted in our paper on the [Changing Roles in Information Governance](#), IG professionals' roles will be forced to evolve, and both IG professionals and lawyers need to continue to stay up to date on the changing landscape of GenAI.

Preservation and auditing

A preservation order, or legal hold, requires a firm to preserve records that are potentially relevant to an anticipated or actual legal action. Going forward, IG professionals will need to consider how to determine if GenAI content is within the scope of such an order.

IG professionals are currently grappling with questions,

similar to the earlier point of 'what is a record?', such as:

- ▶ Does the legal hold apply to records generated by a GenAI tool? Does it make a difference if the output of a GenAI tool was used only as a first draft of the ultimate work product or used without modification?

- Are the original materials from which the content was derived (which may be wholly unrelated to the matter subject to the hold) subject to the hold?
- Does the prompt or prompt sequence used to generate a record need to be preserved, along with the result?
- Given that the output of a prompt can be dynamic, is it possible to preserve when a certain response to a particular prompt may not always generate the same response?
- Similar to some messaging apps, 'conversations' with GenAI tools can be ephemeral. Are there mechanisms to capture the interaction? Is that even worthwhile, given the dynamic nature of responses?
- Can GenAI itself be used to determine which materials are relevant to the implementation of the hold (e.g., by performing a scan or search of data)?

As this space continues to evolve, IG professionals should continue to apply the usual best practices for preservation. Once the duty to preserve is triggered and it becomes evident that GenAI records may be within scope, IG professionals should consider the following:

- Ensure that any GenAI tools you employ have a means of reviewing and validating the audit trail of prompts, inputs, and outputs. The tool should also allow you to capture and preserve associated metadata, such as time stamps.
- Users should maintain records of their own use of such tools, including the specific prompts and inputs, and the responses generated. Such records can help identify content that may be relevant to a preservation order.
- Implement regular backup procedures for prompts and inputs to help ensure their preservation. This is particularly important in the context of dynamic, collaborative tools.

- The preservation procedure should include the issuance of legal hold notices to relevant personnel, including those using GenAI tools, and should specifically emphasize their duty to preserve and prevent inadvertent deletions. If the firm uses a tool that does not allow for a central record of inputs and outputs, be sure to explain an individual's duty to preserve their own prompting records.

Auditing AI systems within a law firm involves ensuring that the deployment and use of AI tools comply with internal policies, client directives, and regulatory requirements. This includes verifying that AI-generated outputs are accurate, reliable, and ethically sound. Review of data inputs, processing methods, and outputs of AI tools should be done to ensure that they do not breach confidentiality or violate intellectual property rights. Additionally, auditing processes should confirm that AI tools are used in accordance with the firm's IG framework, including alignment with or variance from retention and disposition schedules, and that use of client data in AI tools is properly authorized, understood, and documented.

Effective AI auditing also requires continuous monitoring and evaluation of AI systems to identify and mitigate potential risks. This includes assessing the robustness of the AI models, reviewing for biases, and verifying that they operate within the ethical guidelines established by your firm.

Also, ensure that there are mechanisms in place to track and log AI interactions based on your firm's governance requirements. This may include logging of prompts and outputs to maintain transparency and accountability. A strong audit focus helps in maintaining compliance with evolving legal standards and professional guidelines, thereby safeguarding the firm's reputation and client trust.

Matter lifecycle management

Matter and lifecycle management is the process of capturing and maintaining client or matter information, organized by matter type, area of law, or practice area. It includes client engagement documentation and ensures the collection, organization, and access to matter file

content throughout the matter's lifecycle. Additionally, it encompasses systematically closing matters in firm systems upon the conclusion of representation.

GenAI could enhance matter lifecycle management by

automating the capture, organization, and maintenance of client and matter information. GenAI can streamline the onboarding process, reduce the effort to collect accurate documentation, and facilitate efficient access to matter files. Throughout the lifecycle of a matter, GenAI could aid in organizing and retrieving critical information, ultimately improving workflow efficiency and reducing human error. At the conclusion of representation, GenAI could guide compliance efforts to ensure proper archival within firm systems.

GenAI can sometimes produce errors, which could lead to incorrect information being captured or maintained. Because the quality of the data could impact the accuracy of the results produced by GenAI-powered software, focusing on data quality is important. For example, many firms struggle to capture the correct area of law or practice area at the time of matter opening. GenAI could be used to overcome the challenges of manual classification which would have a positive ripple effect on many other IG processes.

Conclusion

The integration of Generative AI into the legal industry brings both opportunities and challenges, particularly in the realm of information governance. While these tools offer the potential to improve efficiency, streamline workflows, and transform the practice of law, they also raise important questions related to privacy, intellectual property, retention policies, and client directives. This paper has explored these implications by examining the IG-specific considerations outlined at a high level in our [August 2023](#) paper and also addressed each IG process component as outlined in the “wheel” from our [LFIGS 2.0 paper: An Established Law Firm Information Governance Framework](#).

A thoughtful approach to GenAI requires transparency with clients, collaboration among key stakeholders, and the development of clear policies to mitigate risks while supporting innovation. By staying informed about legal and regulatory changes and continuously adapting governance practices, firms can effectively manage the complexities associated with this technology. By following the processes outlined in the LFIGS wheel, firms can break down the challenges presented by GenAI into manageable workloads to ensure that as they look to adopt GenAI and capitalize on all its potential benefits, they do so in a way that aligns with information governance requirements.

Footnotes

¹For more information on the legal arguments of copyright owners versus AI platform developers, see [“AI’s Billion-Dollar Copyright Battle Starts with a Font Designer,”](#) by Isaiah Poritz, Bloomberg Law, December 18, 2023.

²[“New Models and Developer Products Announced at DevDay,”](#) November 6, 2023.

³For more detail and analysis, see [“How Copyright Lawsuits Could Kill OpenAI,”](#) by Adam Clark Estes, Vox.

⁴[“Artificial Intelligence \(AI\) Policy Considerations,”](#) by Ryan N. Phelan, Marshall Gerstein, October 31, 2023.

⁵[“AI Regulation is Coming: What is the Likely Outcome?”](#) By Bill Whyman, Center for Strategic & International Studies, October 10, 2023.

⁶[“A Look at AI Regulation in California,”](#) by Louis Lehot and Natasha Allen, Foley & Lardner LLP, January 10, 2024.

Appendix A

What is machine learning (ML)? ML is a part of AI where computers learn to recognize patterns and make decisions from data. It's like giving a computer the ability to learn from experience. In law firms, ML helps lawyers search through legal documents quickly and predict case outcomes by looking at similar past cases.

What is natural language processing (NLP)? NLP is another type of AI that helps computers understand and respond to human language. In law firms, it's used to read and analyze legal documents, find important points, and even help with communication by analyzing the words and tone used.

How do ML and NLP differ from Generative AI?

Generative AI is a newer kind of AI that can create new content, like writing legal documents from scratch. It uses the pattern recognition skills of ML and the language

understanding of NLP to produce new, human-like text. While ML and NLP analyze and interpret existing information, GenAI takes it a step further by generating new information that didn't exist before.

What are large language models (LLMs)? LLMs are advanced GenAI models that are especially good at language tasks. They can write, translate, summarize, and answer questions. They're built using the principles of ML and NLP to understand and generate text similar to what a human would write.

In summary, ML is about learning from data, NLP is about understanding language, and GenAI is about creating new, original content. Together, they're changing the way law firms work by making tasks faster and helping lawyers focus on more complex issues.

Appendix B

A few early cases have emerged:

- **IBM Corp. v. Visentin, 2011 WL 672025 (S.D.N.Y. Feb. 16, 2011):** This case involved a former IBM executive who left the company to join a competitor, Hewlett-Packard, and was sued by IBM for breach of a non-compete agreement. The court granted a preliminary injunction to IBM, finding that the executive had access to confidential and proprietary information about IBM's AI projects and strategies, and that his move to HP would cause irreparable harm to IBM. The court considered the fact that AI was a highly competitive and fast-moving field and that the executive's knowledge and skills were valuable and transferable to HP.
- **Husch Blackwell LLP v. Reeg, 2018 WL 372696 (E.D. Mo. Jan. 11, 2018):** This case involved a former partner of a law firm who left the firm to join another firm and was sued by his former firm for breach of fiduciary duty and tortious interference. The former partner had been involved in a litigation matter that involved AI technology and had allegedly solicited the client to follow him to his new firm. The court denied

the former firm's motion for a preliminary injunction, finding that the former partner did not have access to any confidential or proprietary information about the AI technology and that the client had voluntarily decided to switch firms based on its own assessment of the lawyers' capabilities and fees.

- **Rosenberg v. DVI Receivables XVII, LLC, 835 F.3d 414 (3d Cir. 2016):** This case involved a bankruptcy trustee who sued a law firm for legal malpractice, alleging that the firm failed to advise him of a potential conflict of interest arising from its use of an AI software to analyze the debtor's financial records. The trustee claimed that the AI software was owned by a creditor of the debtor and that the law firm had a duty to disclose this fact and obtain the trustee's consent before using the software. The court affirmed the dismissal of the lawsuit, finding that the trustee failed to show that the AI software was biased or unreliable, or that the law firm had any financial or personal interest in the creditor or the software.



800.899.IRON | [ironmountain.com](https://www.ironmountain.com)

About Iron Mountain

Iron Mountain Incorporated (NYSE: IRM), founded in 1951, is the global leader for storage and information management services. Trusted by more than 225,000 organizations around the world, and with a real estate network of more than 98 million square feet across more than 1,400 facilities in over 60 countries, Iron Mountain stores and protects billions of information assets, including critical business information, highly sensitive data, and cultural and historical artifacts. Providing solutions that include secure storage, information management, digital transformation, secure destruction, as well as data centers, art storage and logistics, and cloud services, Iron Mountain helps organizations lower cost and risk, comply with regulations, recover from disaster, and enable a more digital way of working. Visit www.ironmountain.com for more information.

© 2025 Iron Mountain, Incorporated and/or its affiliates ("Iron Mountain"). All rights reserved. Information herein is proprietary and confidential to Iron Mountain and/or its licensors, does not represent or imply an invitation or offer, and may not be used for competitive analysis or building a competitive product or otherwise reproduced without Iron Mountain's written permission. Iron Mountain does not provide a commitment to any regional or future availability and does not represent an affiliation with or endorsement by any other party. Iron Mountain shall not be liable for any direct, indirect, consequential, punitive, special, or incidental damages arising out of the use or inability to use the information, which is subject to change, provided AS-IS with no representations or warranties with respect to the accuracy or completeness of the information provided or fitness for a particular purpose. "Iron Mountain" is a registered trademark of Iron Mountain in the United States and other countries, and Iron Mountain, the Iron Mountain logo, and combinations thereof, and other marks marked by © or TM are trademarks of Iron Mountain. All other trademarks may be trademarks of their respective owners.