Hogan Lovells

News

Privilege and AI: managing the risks when businesses use AI

31 March 2026

Lydia Savill, Andrew Holland, Reuben Vandercruyssen and Thomas Evans provide practical guidance on how to manage legal professional privilege when in-house legal teams and their non-legal colleagues use generative AI in everyday workflows. This article first appeared in the April 2026 issue of PLC Magazine and can be viewed here.

Introduction

Generative AI (GenAI) is now embedded in day-to-day work across many organisations, involving tasks that may include creating documents, analysing information and drafting supporting communications.

This article is a practical guide on how to seek to manage legal professional privilege when in-house legal teams and their non-legal colleagues use GenAI in everyday workflows. It focuses on the privilege risks in client-side AI use and AI-enabled workflows between the business and the legal

team, including the work intake and triage processes that the legal team operate and lawyer-led review of AI outputs. It does not attempt a standalone analysis of lawyers using large language models (LLMs) to research or draft substantive advice and the distinct issues that can arise.

Many of the AI-enabled workflows discussed in this article engage privilege principles in ways that have not yet been tested in the English courts. Any assessment is therefore likely to be highly fact-sensitive and turn, in particular, on AI tool configuration and confidentiality controls. Against that backdrop, the focus of this article is practical: it offers a simple operating model to help preserve privilege where it can properly arise and to frame the legal arguments that a business can deploy in the future, and explores a four-scenario matrix to put this into context.

Basic concepts of privilege

Legal professional privilege in England and Wales has two key branches: legal advice privilege (LAP) and litigation privilege. This article focuses mainly on LAP.

LAP protects confidential communications between a client and a qualified lawyer for the dominant purpose of giving or obtaining legal advice, including the continuum of communications in the relevant legal context. Importantly, it extends to the whole continuum of communications between the lawyer and their client in the relevant legal context to allow legal advice to be given and received when appropriate. LAP can also protect documents that would reveal the content of legal advice or the advice being sought, even if those documents are not lawyer-client communications. It can also extend to lawyers' working papers and draft client materials that are created to obtain advice, but the circumstances in which privilege protection is available in those situations are not always clear on the current state of the law.

Two pressure points are particularly relevant in the context of generative AI:

- **Confidentiality**. Privilege cannot be asserted unless the document is confidential. While limited disclosure on confidential terms might preserve privilege, any wider dissemination that causes

confidentiality to be lost will be fatal to a privilege claim.

- **The involvement of a lawyer**. A privilege claim is likely to be strongest where the work is clearly part of seeking or giving legal advice, and where a lawyer is actually involved rather than an AI tool operating as a substitute. This interacts with the issue of who the client is. English law takes a restrictive view of which employees count as the advice-seeking "client" for LAP purposes.

The two lanes model

The bedrock of the "two lanes" model mirrors privilege best practices outside of the AI context; that is, work should be split into one of the following two lanes wherever possible:

The business productivity lane. In the business productivity lane, it is assumed that there will be no privilege. AI can be used for drafting, summarising and formatting business-as-usual (BAU) communications, but these should be kept factual and should avoid introducing or seeking legal advice.

The legal lane. For work in the legal lane, approved tools and controlled channels should be used, ensuring tight circulation, clear storage in the legal workspace, and that lawyers are involved at an early stage if it is anticipated that advice will need to be privileged.

No magic switch

There is a persistent misconception around privilege analysis that the mere involvement of the legal team is a "magic switch" that automatically cloaks material with privilege, whether the lawyer is involved from the outset or brought in later on. That binary approach simply does not exist in English law. Privilege is not generated merely because a lawyer is copied in on a communication and nor can it be retroactively found simply by circulating pre-existing commercial material through the legal department.

The correct approach is instead one of function over form; that is, whether the communication was brought into existence for the dominant purpose of seeking or giving legal advice, or its disclosure would reveal the substance of that advice or the legal issues on which that advice was sought.

The idea that merely copying in the legal team on a document or communication will grant privilege is wrong and actively harmful. Businesses should use these guiding principles instead:

- **No backdating.** Circulating already existing material through the legal team at a later date does not make that material privileged.

- **Function over form.** Privilege turns on whether the communication was created for the dominant purpose of seeking or giving legal advice, and whether disclosure would reveal that advice.

- **Circulation kills.** Wide internal sharing can destroy confidentiality even if the legal team is involved (see below).

- **Separate purposes.** It is important to keep business recommendations and legal advice-seeking separate, or clearly signpost which is which.

Confidentiality

Confidentiality is the gateway to any claim to privilege; without it there can be none. Not all AI tools create the same confidentiality and disclosure risk. There is an emerging tendency to treat any interaction with a public chatbot as automatically fatal to confidentiality. That may be too absolute as a matter of principle; under the user agreement, providers of AI tools typically permit the use of inputs for limited purposes such as model improvement, quality assurance or safety review, rather than public dissemination. The more difficult and, as yet, untested question is whether that kind of permitted third-party access or use is enough to defeat confidentiality for privilege purposes in all cases.

This theoretical position should not make legal teams complacent in practice. Recent US authority shows how unforgiving the legal analysis can be where consumer tool terms contemplate provider use and potential onward disclosure. In US v Heppner, the US District Court for the Southern District of New York held that a set of prompt-and-response documents that were created using the AI tool, Claude, and were found on the defendant's devices, were neither attorney-client privileged nor protected work product (2026 WL 436479 (SDNY Feb. 17, 2026)). This was, in part, because the exchanges were with a non-lawyer platform and the applicable privacy terms contemplated the retention and use of inputs and outputs, including potential disclosure to third parties, such as regulators, leaving the user with no realistic basis to treat the conversation as confidential. The court also indicated that even if privileged material had been pasted into the AI tool, sharing it with the provider would defeat that privilege.

The lesson for UK organisations is not that English courts will necessarily adopt the same bright-line test but that confidentiality risk is real and increasingly litigated, so legal work should stay inside approved AI tools and controlled legal processes. Even if there are arguments that confidentiality is not necessarily lost in every scenario, the operational message remains the same: businesses should not take that risk with sensitive material; instead they should use approved, controlled tools and channels for the legal lane.

Before using an AI tool for anything that may be privileged, businesses should ask:

- Where the prompts and outputs live.
- Who can access the prompts and outputs.
- How long the prompts and outputs are retained for and whether they are searchable.
- Whether inputs or outputs are used to train or improve the service.
- Whether outputs are automatically shared or pulled into enterprise search.

The fundamental question is whether the tool creates a risk of third-party supplier access to inputs or

outputs, or an internal access risk through retention, search and broad access, which can undermine confidentiality.

Not all AI is equal

The most common AI tools that are used for legal work fall into four broad categories.

Public chatbots

Public chatbots are consumer-facing applications that provide access to LLMs that are hosted, maintained and operated entirely by the LLM provider. When a user enters a prompt, they are sending content to a third party's systems for processing. The provider may store prompts and outputs, and related metadata, for limited purposes such as service operation, quality assurance, safety review and, in some cases, model improvement under the applicable terms.

Public chatbots should be treated as high risk for privileged information due to the risk to confidentiality. Anything that is typed into a publicly available chatbot should be treated as potentially exposed beyond the organisation, which can undermine confidentiality and make it more challenging to assert privilege later on.

The bottom line is to not enter privileged material or sensitive factual information into public chatbots. A business's internal processes for non-lawyers should reflect this requirement and prohibit the seeking of legal advice through public chatbots. The safest approach is to assume that there is no privilege when using a public chatbot and treat any use as a disclosure-risk decision.

Enterprise copilots

Many organisations now deploy private, enterprise AI assistants, such as Microsoft's Copilot for Microsoft 365, which operate within the organisation's cloud tenant and comply with its identity, access and data governance controls. In practice, these tools typically combine an LLM with enterprise "grounding"; for example, retrieving relevant emails, files and chats from approved sources and passing snippets into the model at run time, and then returning an output to the user.

The key differentiator to a public chatbot is that, if the service is properly configured and contracted, prompts and retrieved context should stay within the organisation's controlled environment and not be used to train a public model. This generally presents a lower confidentiality risk than public chatbots provided that access, logging and retention, and sharing and search settings, are properly controlled. How the provider handles data is also key; even a supposedly private model might use inputs to improve its service, so the contractual position in that respect should be confirmed.

This does not mean that the use of a closed or proprietary AI tool will somehow allow for privilege protection, as confidentiality is just the gateway.

If non-lawyers use an internal AI assistant to generate analyses or draft reports on legal issues without involving lawyers, those outputs will still face the same challenges for privilege due to the lack of lawyer-client communication. However, these systems represent a better approach from the confidentiality perspective, provided that the tools in question are carefully vetted and approved.

These tools can be viable for the legal lane from a confidentiality perspective if access, retention, searchability and training settings are controlled. Risks do not come only from the supplier; internal access and inadvertent circulation also present risks.

Specialist legal tools

The market increasingly offers AI tools that are built specifically for legal teams, ranging from e-disclosure platforms with AI capabilities to lawyer-facing assistants that are integrated into document and matter workflows. These tools are typically designed for controlled environments; that is, matter-based workspaces, with granular permissions, audit trails and deployment models that sit behind organisational security controls. This "legal-grade" design does not make them risk-free from a confidentiality or privilege perspective. The outcome still turns on the tool's characteristics, the contractual position and how it is used in day-to-day operations.

They are often easier to govern than general consumer tools because they can be ring-fenced by matter, access can be restricted and usage tracked. However, it is still important to understand the basics around retention, the use of inputs to improve models, and who can access content, including supplier support and administrators. Where the tool connects into document stores or email, it is worth checking what gets indexed or cached, and whether access controls and ethical walls carry through cleanly.

Agentic tools

AI tools that operate as a fully autonomous or semi-autonomous agent inside an organisation's workflow are attracting increasing attention. They can be seen as more like a junior team member than a drafting tool. Depending on how it is configured, an agent may be integrated across multiple systems and able to take actions, as well as generate text, that span both legal and non-legal domains.

Because these systems are more prone to blurring legal and operational functions, they can be harder to govern from a privilege perspective. An agent that is supporting project management might pull together information from across the business and produce recommendations that mix operational views with legal analysis. That output may then be shared beyond the legal team, making it harder to maintain confidentiality and to show that the work was done for the purpose of legal advice or litigation. Where agents are connected into document management systems, ticketing tools or email, it is also worth understanding what they can access, what they can write back, and what logs are retained.

Operational basics

Privilege is not a product feature; it is preserved or lost through confidentiality, purpose and controlled circulation.

The controls set out below are the baseline requirements that need to be satisfied before getting into any detailed factual analysis. They are designed to prevent two common privilege problems: losing confidentiality and blurring legal work with business work. To help to prevent these problems,

businesses should:

- Check the AI supplier terms, in particular, to confirm that the supplier will not train its AI tools on the business's inputs by default.
- For legal work, use tools only from an approved tool list that sets out named approved tools and prohibited tools.
- Use matter-based storage for AI artefacts and outputs, only retain what is necessary and restrict access to privileged work product only to those who need to know (see "AI artefacts" below).
- Understand the retention and search settings, and assume that logs are part of the record.
- Restrict the use of connectors, such as document management systems, email platforms and SharePoint, as far as possible when using any AI tool.
- Share the minimum information necessary for the task at hand, be structured about the use of data, and be intentional about what context is provided to an AI tool.
- For anything that the business may later want to argue is privileged, keep workflows lawyer-led, under legal direction and supervision, and ensure that they are undertaken for the purpose of legal advice or litigation.
- Avoid using AI as a general business adviser on the same thread as legal advice. Legal analysis should be kept separate from BAU tasks.

An AI and privilege workflow

Before using an AI tool on a live matter, users in a business should consider whether:

- The matter is sensitive. If it is, users should make sure that the AI is an approved tool, use matter-based storage, and only share the minimum information necessary for the task at hand.

- The output will be shared outside of the legal team. If it will, users should pause to consider whether they are comfortable with that.

- The tool is a connected tool or an agentic tool; that is, an AI system that can do more than generate text and can take actions in other software. If it is, users should check connector scope, logs and writeback controls; that is, the safeguards that govern whether (and, if so, how) the AI tool can effect changes to the business's operational systems, instead of only reading data and producing an answer.

Two red lines

There are two important rules:

- Do not enter privileged material into unapproved tools. Users in the business should receive training on what constitutes privileged material, including draft advice, chronologies prepared for counsel, advice emails, investigation notes and meeting minutes.

- Escalate a matter where appropriate. If the question is bespoke or high stakes, the matter should be escalated to a human lawyer rather than iterated internally with AI.

AI artefacts

AI tools tend to generate and retain a wider set of artefacts than traditional workflows, including prompts and retrieved context, iterative chat threads, intermediate drafts, and (in meeting contexts) verbatim transcripts, summaries and action lists. These records can be shared automatically, retained for long periods, or indexed into enterprise search or e-discovery, which increases discoverability and widens access, even in situations where this was unintended.

The practical consequence is straightforward: prompts, chat history, intermediate drafts and system logs should be treated as part of the record and managed in the same way as any other sensitive working papers because if they are not privileged or do not reveal privileged advice, they are potentially disclosable.

Four scenarios

Below are four typical scenarios involving AI-generated documents in a corporate setting, which consider how the principles discussed above may be engaged.

Internal-only legal chatbot

In the first scenario, employees use an internal-only legal Q&A AI chatbot to raise questions and receive AI-generated answers, with no lawyers in the chain of seeking or providing those responses.

In the most basic example of this type of scenario, in which an employee asks a general-purpose chatbot for legal advice, the base case is that the prompts and the outputs would not be privileged. LAP is confined to confidential communications between a lawyer and their client for the dominant purpose of giving or obtaining legal advice. That structure is not obviously present where the exchange is solely between a non-lawyer and an AI tool.

Best framing for privilege. While still uncertain and untested, a best-case factual pattern for this broad legal Q&A chatbot scenario might be a tool that is narrowly directed at being a delivery channel for lawyer-written advice, or a structured triage step into the legal team. For example, consider a chatbot that:

- Under the guardrails to maintain confidentiality discussed above, is trained on a curated internal knowledge base of existing, lawyer-drafted material with clear version control and approval.

- Is overseen by the legal team to ensure that the tool is only retrieving, summarising or condensing lawyer approved guidance and is not generating novel advice, with periodic review of queries and outputs to confirm that this remains the case.

- Is accessible only by a defined set of authorised client users; that is, employees who are tasked with seeking and receiving legal advice on behalf of the organisation. In addition, there is clear written guidance for its use, including that any recipients may only use it for defined internal

purposes and must not disseminate it more widely, to avoid undermining confidentiality and waiving privilege.

- Escalates bespoke legal queries to a human lawyer, with the escalation and any subsequent lawyer advice captured within the normal lawyer-client channel.

Where these features apply, the tool in question will look more like the retrieval and display of existing lawyer-drafted guidance or a mechanism to route requests into the legal team, rather than the generation of new advice outside of the lawyer-client relationship.

Potential challenges. The biggest weakness in this scenario, even on the best case identified, remains the absence of a lawyer in any exchange with the AI tool that is not escalated. The risk is that the initial query and output will be seen as the non-lawyer seeking guidance from the tool instead of seeking legal advice from a lawyer, rather than the delivery channel interpretation. It will also be apparent that, in practice, tailoring an AI tool to this kind of delivery channel model may require additional design and governance; for example, a tightly curated knowledge base, defined user groups and clear escalation to a lawyer. Some organisations may conclude that this is more restrictive than they want for general use.

There are also two practical vulnerabilities. Firstly, where the AI tool is used for mixed, part-legal and part-operational, helpdesk queries, it may be difficult to show that any given prompt or output was created for the dominant purpose of giving or obtaining legal advice. Secondly, if access is broad, or not clearly limited and documented, many interactions may be treated as communications by non-client employees, weakening any LAP claim even if the legal team is responsible for the tool.

Practical tips. A key factor in presenting the best argument for preserving privilege in this scenario may be to use tools with clear dividing lines between:

- A business-use lane, which deals with general business help and non-legal information, where no legal advisory queries should be handled and there is no expectation of preserving any privilege.

- A legal-use lane where any legal advisory queries should be handled using secure tools and human lawyers (see *"The two lanes model"* above).

In practice, these Q&A chatbots fall on a spectrum. At one end, the chatbot is purely a business self-service tool where users ask questions on legal issues and receive AI-generated answers with no involvement from the legal team. Based on current law, this type of chatbot is the least defensible from the perspective of LAP. At the other end of the spectrum is a retrieval-first delivery channel that primarily finds lawyer-approved guidance from a curated knowledge base, which is shared on a controlled limited waiver basis, escalates bespoke questions, and follows an AI-assisted intake and triage process by the legal team where the user's query is routed into the legal team and a lawyer reviews and responds. Although untested, there is more scope for arguments that privilege should be maintained in this type of chatbot. However, even this is not a silver bullet and there remains a real risk of waiver of privilege over the pre-existing advice without tight guardrails.

For chatbots that sit somewhere in between, such as a chatbot that generates responses drawing on legal inputs but with limited oversight by the legal team, the privilege position is more fragile still. There is a real risk that the advice generated, and the previous advice on which it is based, will lose any claim to privilege. The risk can be reduced, but not eliminated, by the use of guardrails, such as clear escalation to the legal team, restricted access, clear conditions on which responses based on existing advice are shared, controlled storage and circulation, and retention and logging controls. On the current state of the law, caution is warranted.

AI draft created but not sent

In the second scenario, a business team uses an AI tool to draft a briefing, such as a collation of facts and a list of concerns and questions, which they intend to use to brief and seek advice from the legal department. Despite that intention, they never actually do so.

Best framing for privilege. The intention to seek advice at some point is not, by itself, enough to attract privilege. The question is whether the draft was brought into existence for the dominant purpose of obtaining legal advice. There is no clear general rule that a draft must have been

transmitted in order to attract LAP, and there are good policy reasons for protecting draft requests for advice that are prepared for the legal team within a genuine legal intake process. The most defensible framing is therefore narrow and fact-sensitive: where a draft is prepared by or for the defined advice-seeking client group, for the dominant purpose of seeking legal advice, and it is kept confidential in the legal lane, there is a principled argument that disclosure would reveal the advice sought and the relevant legal context, even if the draft was not ultimately sent.

Potential challenges. While there may be a defensible privilege argument depending on the facts, three practical hurdles need to be overcome:

- Satisfying the dominant purpose test; that is, whether the draft was created mainly so that the legal team could advise on it, rather than as a BAU analysis that later became relevant.

- Showing who the client group is. For the purpose of ensuring that privilege applies, it matters whether the draft was produced by, or for, the defined group that is tasked with seeking and receiving legal advice on behalf of the organisation.

- Demonstrating confidentiality. Wide sharing, messy storage and uncontrolled logs can make it harder to maintain confidentiality and to present the draft as part of a controlled advice-seeking process. Where legal and non-legal content are mixed, it may be possible to separate or redact information, but the dominant purpose and confidentiality are still the critical organising principles.

Practical tips. At its most straightforward, if a document is genuinely intended to be used as part of seeking legal advice, it should be treated as a draft communication to the lawyer, and it would be better to send it. The other usual recommendations apply around restricting access and circulation to preserve confidentiality, and ensuring that any such documents do not also deal with business-type enquiries.

Where an organisation has set up a legal lane; that is, a genuine intake and triage process for the legal team to deal with legal queries, and users are preparing a draft specifically for the legal team within that controlled channel, there are good principled and policy reasons why that draft should be treated as part of the advice-seeking process and protected accordingly. In practical terms, this is best supported by using a clear legal-intake route; for example, a portal where users submit or save draft queries to the legal team, limiting access to the defined advice-seeking group and the legal team, and storing the draft within the matter or legal workspace so that it sits within the ordinary lawyer-client workflow.

Instructions sent to legal team

The third scenario is a development on the second scenario, in which a business team again uses an AI tool to draft a briefing to the legal team but, this time, actually sends it.

Best framing for privilege. Where a communication with a lawyer takes place, it will be subject to a more conventional privilege analysis. In practice, it will usually be necessary to evidence two points:

- That the draft was created and sent for the dominant purpose of obtaining legal advice.
- That it sits within a controlled advice-seeking workflow, with confidentiality preserved, and circulation limited to the defined advice-seeking group and the legal team.

In addition, the established principles around the continuum of communications mean that a court should, in principle, consider the material that forms part of the continuum of the advice sought by the client to be part of the privileged communication, provided, as ever, that confidentiality is maintained.

Potential challenges. Even though this scenario fits more straightforwardly within the existing legal principles, there may still be potential challenges.

Firstly, care should still be taken as to the confidentiality of the AI-drafted content before circulation to the lawyer; wide pre-circulation should be avoided. Secondly, there is still a need to ensure that the document has been created for the dominant purpose of seeking legal advice. An AI-generated document in this scenario that mixes questions of business strategy and legal advice is likely to cause challenges. This will often be highly fact-sensitive; for example, using AI mainly to triage whether to approach a lawyer, or to develop an internal strategy to discuss with a lawyer later, may fall short of the dominant purpose test.

Where possible, the draft and any iterations should be kept inside the legal lane, adhering to an approved tool, matter workspace and need-to-know circulation, and avoiding sending multi-addressee versions around the business "for views" before the legal team has been instructed. Further, the wider AI artefacts should be treated as part of the record: earlier internal iterations, the AI chat and prompt history that generated the draft, and any automatically captured logs can become the weak point if they are widely accessible, retained indefinitely or held by a third-party provider.

Practical tips. This scenario largely follows familiar privilege principles in relation to keeping the draft confidential, focused on obtaining legal advice and managed within the legal lane.

Non-legal briefing later turned into instructions

In the fourth scenario, a business team generates an internal document outside of the legal advice context that is not initially intended to be provided to the legal team. The business team then decides to seek legal advice and the document, or parts of it, is sent on to the legal team as part of that request for advice.

Best framing for privilege. Existing principles make clear that LAP is not created retroactively. A document that is generated outside of a legal advice context does not become privileged simply because it is later sent to a lawyer. The better way to think about this scenario is that the original document usually keeps its original, often non-privileged, disclosable status, but new lawyer-client communications around it can be privileged, and records or summaries that evidence or paraphrase

the substance of legal advice may attract LAP. This privileged layer of communications, that is, the request for advice, the advice and any lawyer-led summaries that reveal the advice sought or given, should be consciously created and, provided that confidentiality is maintained, should be protected.

Potential challenges. The core point is straightforward: a pre-existing business document does not become privileged just because it is subsequently sent to the legal team or referred to when seeking advice. Therefore, the analysis is not whether the original document can become privileged, as it generally cannot, but whether the privileged wrapper around it can be protected and preserved.

Three things will make the position harder in practice:

- Confidentiality. If the original document has already been widely shared, it may be more difficult to keep any advice-related layer of communications confidential.

- Mixed-purpose communications. If communications are as much about operational decision making as legal advice, the dominant purpose point becomes contested.

- Document hygiene. If the business recirculates lawyer advice or blended summaries in uncontrolled channels, it becomes harder to draw a clean privilege line.

Practical tips. The best approach is to avoid landing in this scenario in the first place by involving the legal team at an early stage and keeping early work in the legal lane.

A privileged layer around the pre-existing material should be deliberately created and protected. In practice that means:

- Sending a clear, confidential cover communication that frames the request as one for legal advice so that the request sits within LAP.

- Extracting and sending only what the legal team needs, rather than forwarding a widely-circulated pack wholesale and, where helpful, having the legal team prepare a new lawyer-led summary or issues note that evidences the advice sought and given, and forms part of the

continuum of seeking and receiving advice.

The practical message is not to send pre-existing documents through the legal team in the vain hope that this will grant them privileged status; instead, a better approach is to involve the legal team early, keep circulation tight, and focus on creating a clean, confidential advice-seeking channel and a lawyer-led record of the advice process.

Litigation privilege

The second branch of legal professional privilege under English law, litigation privilege (LP), can be extremely helpful because it is not limited to lawyer-client communications. Where its conditions are met, LP can protect confidential communications between a client and a lawyer, and between either of them and third parties, that are created for the dominant purpose of conducting, avoiding or settling adversarial proceedings that are pending or reasonably in contemplation.

The three gateways to clear for litigation privilege (LP) to apply are:

Adversarial proceedings. LP protects material that was created for adversarial proceedings, not purely inquisitorial or fact-finding processes. It matters whether proceedings are genuinely in prospect, rather than there being a general possibility that proceedings might one day follow.

Reasonable contemplation. Adversarial proceedings must be a real likelihood, not a mere possibility or general apprehension. In internal investigations, the fact that it is not yet known what the facts will show, or the fact that there remains some conditionality as to whether there will, in fact, be adversarial proceedings, does not necessarily prevent proceedings being in reasonable contemplation but it should be possible to explain what changed, when, and why it was concluded that proceedings were realistically in prospect.

Dominant purpose. LP requires that each relevant communication was created for the dominant purpose of conducting the litigation. The courts test this on a document-by-document basis. Dual-purpose materials are a common failure point since, if the same AI tool or workspace is used for business-as-usual problem-solving and litigation preparation, it becomes harder to prove dominance.

However, LP is not a general-purpose safety net. Courts apply the tests closely and will look for contemporaneous evidence, not labels applied after the fact.

LP does not convert a pre-existing document that was created before litigation was reasonably contemplated into a privileged document. Feeding a non-privileged business document into an AI tool for litigation analysis does not make the underlying document privileged. The focus instead should be placed on protecting the litigation preparation materials that are generated around it, such as lawyer instructions, evidential analyses and strategy notes, and being ready to show when the LP trigger was met.

Evidence needed if LP is later challenged

If it may become necessary to rely on LP at a later stage, it is advisable to build an evidential trail while the work is in progress. This could include, for example:

- A dated note or a matter-opening record explaining why proceedings are reasonably contemplated.
- An investigation plan or instruction email showing the dominant purpose.
- Clear separation between litigation workstreams and BAU work.
- A record of tool configuration, such as retention, logging and access controls, showing confidentiality control.

In the context of AI, the operational consequences for businesses are to:

- Segregate litigation workflows, such as tools, workspaces and channels, from BAU.
- Treat prompts, outputs, chat history and logs as part of the record, with controls over access, retention and searchability.
- Pay attention to supplier controls, avoiding tools and terms that allow broad supplier use of the business's data, and for LP-sensitive work, ensuring that confidentiality protections are contractual and technical, and that access is tightly limited.

Meeting notes and AI transcription

The automated transcription and summarisation of meetings is a common use of GenAI tools. Meeting minutes are not automatically privileged simply because a lawyer attended.

LAP turns on the dominant purpose of the relevant communication or document and whether it is confidential. In the meeting context, this usually means that minutes or transcripts are privileged only to the extent that they record, paraphrase or would reveal the substance of legal advice sought or given, or the necessary exchange of information whose object is the giving of legal advice as and when appropriate. Factual records of what happened, business updates and decisions, or action lists will often be non-privileged unless they would, directly or by necessary implication, disclose the legal advice.

This becomes especially important for mixed meetings. Even if a member of the legal team is present, a mixed commercial and legal meeting does not become wholly privileged. Where the content is severable, privilege may apply only to the legal parts and the rest may be disclosable. In a corporate setting, the "who is the client" point can also be relevant because, if minutes capture communications from employees outside of the defined advice-seeking group, those parts may be harder to bring within LAP even if directed to a lawyer. This is one reason why it can be particularly helpful if LP is engaged, because LP can, where its tests are met, protect a wider range of meeting communications and records, including communications with third parties that are created for the dominant purpose of litigation.

AI tools magnify these risks because they tend to capture and distribute more than a human minute-taker would. Verbatim transcripts record everything, both privileged and non-privileged, automated summaries and action lists can pull legal advice into a non-privileged format, and default settings can auto-share outputs to all attendees or index them into enterprise search.

If privilege may apply to the AI transcription that is used to record a meeting, businesses should:

- Decide upfront whether the meeting will cover legal advice, business-only matters or mixed topics.
- Use the right record for the job. Verbatim transcription should be avoided for mixed meetings unless there is a clear plan for review and redaction; a lawyer-owned note is to be preferred for the legal segments.
- Ask the legal team to review, correct and finalise any transcript or summary before it is circulated.
- Where feasible, separate and clearly label the parts that record the legal advice that is sought or given from factual or business discussions and actions.
- Turn off default auto-sharing, circulate the transcript on a need-to-know basis only, and store the privileged record in the legal workspace or matter file.
- Manage the artefacts by treating prompts, chat history, draft summaries and system logs as part of the record, know what is retained, who can access it and whether it is searchable.

Key Takeaways

AI does not change the privilege rules; it changes the document trail. In practice, the problem is that AI makes it easy for businesses to produce neat-looking notes, summaries and positions on legal matters before the legal team is involved, then share them around the business and save them in

places that are searchable later. AI forces privilege rules that were built around human communications to be applied to novel workflows, so edge cases and grey areas are inevitable. All is not lost. If BAU tasks are separated from legal work, and legal queries are kept inside a controlled, lawyer-led process using approved tools and matter workspaces, keeping circulation to a need-to-know basis, this will give businesses the best chance of preserving privilege.

Six golden rules

The following steps will help to give the business the best chance of preserving privilege.

Pick the correct lane at the start. Teams should decide whether work should go in the business productivity lane, which assumes that no privilege applies, or the legal work lane, in which privilege is preserved as far as possible.

Route legal questions through the legal team. Teams should use an approved intake process and escalate high-stakes issues to a lawyer.

Use approved tools only when in the legal lane. Teams should ensure that the approved tools are configured properly with limited access, controlled logging and retention, and the restricted use of connectors.

Keep it confidential and on a need-to-know basis. Businesses should restrict access to the advice-seeking group and the legal team, and avoid broad circulation and auto-sharing.

Keep it matter-based and tidy. Teams should store and retain only what is necessary in the matter workspace, and avoid scattering outputs across chats and general document or file management systems.

Assume that prompts, drafts and logs are part of the record. Teams should manage retention and searchability, and not generate material that the business would not wish to defend in disclosure.

This article first appeared in the April 2026 issue of PLC Magazine and can be viewed here.

Authored by Lydia Savill, Andrew Holland, Reuben Vandercruyssen, and Thomas Evans.

Contacts

Lydia Savill

Partner

London

Email me

Andrew Holland

Counsel

London

Email me

Reuben Vandercruyssen

Senior Associate

London

Email me

Thomas Evans

Associate

London

Email me