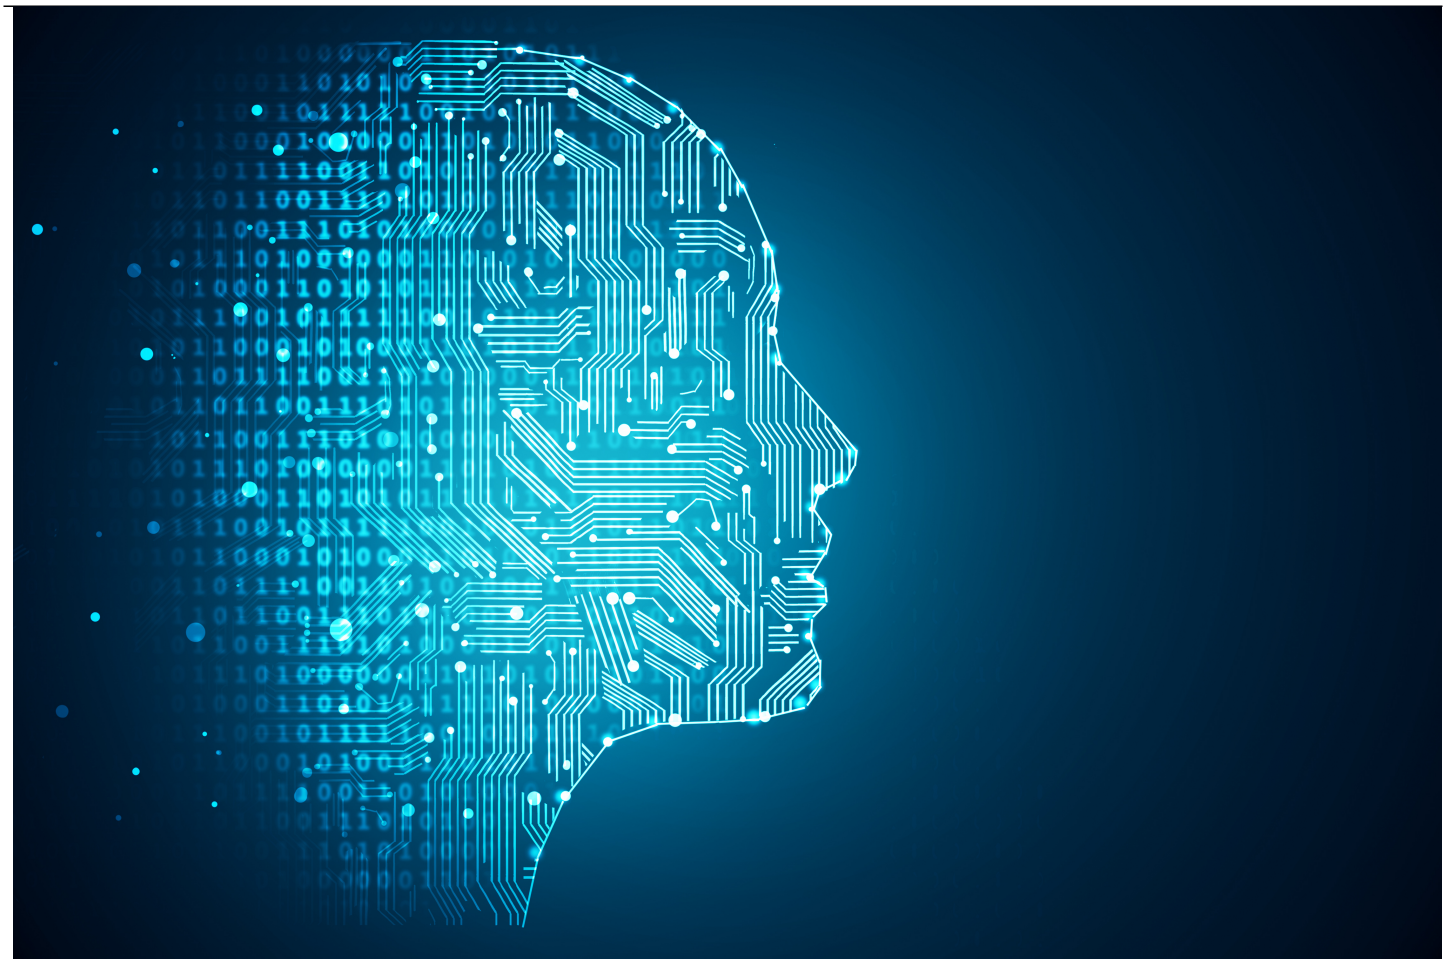


AI Counsel: Clean Data Is the Next Frontier — Why Personal Information Discipline Now Determines Your AI Leverage

Technology, Privacy, and eCommerce



Banner artwork by Peshkova / *Shutterstock.com*

Most conversations about enterprise AI adoption still start with the technical solution: which tool, which model, which vendor. Consider a different starting point. When you graduate from one-off prompts to using AI to help with real work across your emails, meetings, notes, and documents, the single biggest lever on how useful your AI assistant is turns out not to be the tool at all. It is the quality, currency, and organization of the data sitting within the reach of your own keyboard. Increasingly, it lives in the notes you did or did not take.

Instead of visiting a standalone prompt window with a question, your prompts can sound more like this:

- **Microsoft Copilot**: Search my full Microsoft 365 tenant and answer from the available records only. Flag gaps, conflicts, and source limits. Draft a model email to my CEO reconstructing the history, identifying the current status, and explaining next steps.

- **Claude for Work / Claude Enterprise**: Search all connected company sources you have access to, including Google Drive, Slack, email, calendar, shared files, projects, and any enabled connectors. Answer only from available source material, and clearly flag gaps,

conflicts, or limits in access. Generate a single slide providing a dashboard of key department priorities, status, and key actions through the end of the year.

- **Google Gemini / Google Workspace**: Search across my available Google Workspace content, including Gmail, Drive, Docs, Sheets, Slides, Calendar, Meet notes, and shared files. Answer only from what you can verify in those sources, and clearly flag gaps, conflicts, or access limits. Summarize the key developments in this litigation, along with key dates and upcoming deadlines.

The retrieval problem behind enterprise AI

Copilot, Gemini, and Claude differ in interface, ecosystem, and commercial model, but for enterprise users they depend on the same basic architecture: a language model is connected to the organization’s own stored work product through approved systems and connectors.

- In Microsoft, that means content available through Microsoft 365;
- In Google, it means Workspace sources such as Gmail, Drive, Docs, Calendar, and Meet; and
- In Claude’s enterprise environment, it may mean connected repositories such as Google Drive, Slack, project knowledge bases, or other approved systems.

The practical point for lawyers is simple: these tools do not “know” the business because they are smart. They become useful when they can retrieve reliable source material, such as meeting notes, emails, drafts, decisions, redlines, and follow-ups. If those materials are current, well-labeled, and saved where the tool can access them, the assistant can summarize, compare, draft, and reason over them. If they are missing, scattered, stale, or locked away in places the tool cannot reach, the result is either silence, shallow generalities, or confident-sounding reconstruction that should not be

trusted.

Two layers of data governance, not one

The instinctive response to any data quality conversation is to reach for organizational solutions: a central contract repository, clause library, records retention policy, or matter management system with disciplined metadata. Those tools remain necessary, and AI will not compensate for their absence. But they are only half the picture.

The other half is personal. A lawyer's calendar, inbox, meeting notes, drafts, and saved documents are also a data set that AI draws on constantly, whether or not we think of them that way. If you want an enterprise AI tool to draft a credible board update, summarize a negotiation, or flag a shift in a counterparty's position, those everyday artifacts need to exist in usable, searchable form. If they are missing, scattered, stale, or buried in random notes, the best enterprise data governance program upstream will not save you.



The memo to file, rediscovered

Lawyers who trained at firms will recognize this discipline, because we used to practice it as a matter of professional survival rather than AI enablement. The memo to file, the file note after a client call, the contemporaneous record of a partner's oral instruction or a client's verbal waiver, existed because memory is unreliable and because a clean written record protected the lawyer, the client, and the matter. Many of us let that habit atrophy once we moved in-house, where the pace is faster, the meetings are more numerous, and the incentive to write things down for their own sake felt weaker.

AI retrieval tools have quietly restored the incentive. A short, well-structured note after a meeting is no longer just professional hygiene; it is the raw material the assistant needs to produce a good executive summary next week, or to correctly reconstruct the history of a negotiation eight months from now when someone asks how a particular term was resolved. The habit lawyers were trained into at firms turns out to be exactly the habit that makes an in-house AI assistant genuinely useful.

What this looks like in practice

Close the loop on every meeting. A few lines in Teams, OneNote, or a saved email; decision, rationale, owner, next step, captured within the hour, not reconstructed from memory days later.

Write titles and headers for retrieval, not just for yourself. Semantic search performs far better against a note titled "Acme MSA — liability cap negotiation, 8/14" than one titled "notes 3." And for searching through prompts in your AI tool, title each query for easy retrieval and iteration.

Treat verbal guidance the way firm lawyers treated partner instructions. A short confirming email or file note after informal guidance from the GGC or a business leader creates a record the assistant can find later, and creates the same protective record it always did.

Curate, don't just accumulate. A raw meeting transcript is not the same as a clean note. Five minutes spent trimming a transcript or dictated recap into a short, accurate summary is worth far more to future retrieval than an unedited wall of text.

Keep drafts and redlines where the assistant can see them. A contract version saved to a personal desktop folder or a local downloads directory is invisible to tenant-wide retrieval in a way that the same file saved to OneDrive, SharePoint, or the equivalent shared space is not.

Ease into your next-level AI usage

The payoff is ease. A lawyer who keeps clean notes, clear file names, accessible drafts, and current decision records will experience AI less like a search and trivia tool and more like a well-briefed colleague. The work becomes lighter: fewer inbox scavenger hunts, fewer reconstructed timelines, and more time spent applying judgment where it matters.

Join ACC for more AI insights!

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

Spiwe L. Jefferson



Division General Counsel, Division General Counsel, Electrification Service

ABB

Spiwe Jefferson leads global legal and integrity functions across international operations for ABB's Electrification Service division. A sought-after speaker and thought leader on AI enterprise adoption, she has achieved over 710 hours (18 weeks) of AI-driven efficiency gains on one platform alone. Spiwe authors the *Mindful in 5* book series and podcast, providing leaders with actionable strategies for resilience and growth.