




The GC's New Mandate: Architecting AI Governance in Real Time

Technology, Privacy, and eCommerce



Banner artwork by [tete_escape / Shutterstock.com](#)

In February, a federal judge in the Southern District of New York [ruled](#) that a defendant's written exchanges with a public generative AI tool were not protected by attorney-client privilege or the work-product doctrine. The reasoning: the model is not a lawyer, the prompts were not confidential, and material does not become privileged simply because it is later handed to counsel. A settled principle was simply applied to a tool that millions of employees now open without a second thought. That is the gap every General Counsel is now operating in, and it is where the job is being redefined.

If your company is like mine, AI has become a near-obsession. How should we use it to reimagine parts of our business, gain operational efficiency, code faster, and serve our customers with new experiences that help them work better, faster, and smarter? Teams across the company are answering these questions in creative ways, often using the tools closest at hand. The privilege ruling is a reminder that the same speed creating all this value is also creating exposure, and someone has to own both at once. Increasingly, that someone is the GC.

The privilege ruling is a reminder that the same speed creating all this value is also creating exposure, and someone has to own both at once. Increasingly, that someone is the GC.

[BCG's research](#) underscores that AI-innovative businesses thrive when they overcome new barriers. What is not always top of mind inside companies today is how AI reshapes the legal team itself and the barriers it raises around privacy, compliance, and regulation. The privilege ruling is one example. Shadow AI changes the game. Across the private sector, employees are routinely pasting contracts, customer data, and source code into consumer tools that were never reviewed, and that data can be retained, used for training, or disclosed in ways the company cannot claw back. GCs are faced with the question of whether they should guide, and even embrace, AI in the legal department and across the company while the rules are still being written.



The rise of AI coincides with the ascent of the GC to a comprehensive business strategy partner for the entire company. Legal departments just a few years ago were generally focused on expected tasks such as contract reviews and the occasional risk or compliance endeavor. But with AI contributing to every part of the organization, a contemporary GC must reexamine priorities and be prepared to enable different parts of the business.

The rise of AI coincides with the ascent of the GC to a comprehensive business strategy partner for the entire company.

Protecting our IP from absorption into LLMs is a baseline priority. Ensuring our employees do not inadvertently infringe copyright laws across various countries is another key concern. Ensuring code is developed or licensed appropriately is of heightened importance in the age of AI. And of course, legal teams use AI to support legal research, whether it is contract or employment law across jurisdictions, legal trends, new laws that take effect regularly, or new compliance standards that impact our role.

Ensuring code is developed or licensed appropriately is of heightened importance in the age of AI.

What excites me is how AI is changing the GC role in more ambitious ways. While my team and I manage the daily opportunities and risks of AI, we are also building the foundational processes and protocols our business will run on. That is a shift in posture. The GC is no longer only the person who reviews what others build. The GC is the architect of how the company adopts AI, and the source of the guardrails that let it move fast without driving off the road. In my own seat, I have come to think of the role as part architect and part chief change officer, because writing the framework is only half the work. The other half is moving people through it.

Governance in real time

The instinct for any lawyer is to get ahead of a problem: scope the risk, write the policy, then let the business proceed. AI has taken that luxury away. The technology changes weekly, legislation is fragmented and trails adoption and use (often by years), and there is rarely time to come to legal before something happens. The honest reframe is that we no longer have the runway to govern ahead of time. We have to govern in real time. Governance has to be embedded in the workflow and built to change tomorrow, because tomorrow it will need to.

The honest reframe is that we no longer have the runway to govern ahead of time. We have to govern in real time. Governance has to be embedded in the workflow and built to change tomorrow, because tomorrow it will need to.

That is why I believe you cannot govern AI well without practicing it. The risks of a tool become tangible only once you have used it, and you cannot design a credible framework from the outside. So, my legal team adopted AI ourselves, ahead of the rest of the business, rather than behind it, with a high bar for use and real investment behind it. We've had 100 percent adoption across the team, which is the opposite of how legal departments usually behave. We are rarely the first movers. Choosing to be here was deliberate. Doing the work firsthand is what lets us see where words and processes quietly impede adoption, and it gives us empathy for every other team being asked to change how they work. A leader who has actually used these tools understands immediately why a prompt typed into a consumer account is no different from a message that can be subpoenaed, and can write guardrails that reflect that reality.

A shift to continuous compliance

We are moving rapidly from an era of fixed compliance toward a model of constant scenario planning. Global regulations right now are fragmented and evolving. The [EU AI Act](#) scopes some boundaries and various US states have early directional guidance, while the federal government and states still work out who has jurisdiction, and many countries have not taken up AI regulation at all. None of that is a reason to wait. AI is a new technology, but most of what it touches is already governed. Existing principles around confidentiality, privilege, IP, and data protection applied to the fax machine, then to email, then to the cloud, and they apply to AI now. The February ruling applied an old doctrine. Our codes of conduct do not change. We adopt, then we adjust. My instinct, drawn from years in FinTech, is to educate regulators about new technology rather than hide behind the absence of rules.

Looking out to the remainder of 2026, my advice is to stop waiting for a definitive rulebook, because we don't know what is coming on that timeline. AI will keep developing in a historic and difficult-to-anticipate trajectory, and regulation will continue to lag. When formal rules do arrive, they will often be outdated on contact. A clear near-term signal will emerge from the courts. The privilege ruling is an early example of how existing doctrine is being applied to AI, and more privilege, discovery, and liability questions will surface in litigation well before any statute settles them. We have already seen attorneys sanctioned for filing unreviewed AI-generated briefs, and that category of cases will grow. The practical move is to treat case law as the leading indicator, watch it closely, and keep governance flexible enough to absorb each development as it lands. Successful GCs in this era will be nimble, will have gamed out scenarios in advance, and will keep their boards, employees, and customers equipped for change rather than surprised by it.

Successful GCs in this era will be nimble, will have gamed out scenarios in advance, and will keep their boards, employees, and customers equipped for change rather than surprised by it.

The architect's mandate



In a high-velocity business, risk is an accelerant for the goals we are trying to reach, provided you can see it clearly and price it honestly. I have never treated risk as something to simply be eliminated. This is even more relevant given the reality of AI's pace. AI sharpens both the upside and the exposure, and the GC who can hold those together is the one who lets the business move. Businesses should maintain their commitment to control, security, and privacy throughout. [Forrester reports](#) that AI adoption should emphasize security, quality, and privacy as core design principles, and that mindset starts with leadership, especially the GC, who must set the tone and focus priorities.

A practical framework to move fast

Where to focus your strategy with such fast-moving change can be challenging; here are some foundational steps that should never be ignored.

- Approved tool architecture: define which AI tools may be used for what categories of data and use cases. Use the stoplight approach, with a focus on broadening the categories of green and orange/amber to allow the business to innovate and move at its desired pace.
- Privilege and litigation protocol: prohibit employees from using public AI tools to analyze disputes, investigations, or legal exposure unless directed by counsel.
- Continuous review loop: update AI governance based on litigation, regulatory developments, incidents, new use cases, and lessons learned from day-to-day use. Create an AI Forum/Council comprising of stakeholders from Legal, Security, Product, Engineering, Data, etc.
- Trust and adoption: The ability to have oversight of AI outcomes earns trust, not just with the legal team but also employees. In Tipalti's [State of AI in Finance](#) research on AI trust,

professionals said the ability to review AI actions (55 percent) and custom-configure workflows (55 percent) are the capabilities that matter most. Your team and the employees you are building for will move fastest when experts have made the systems around them understandable and accountable.

Also, it goes without saying that every GC today should embrace AI for their own legal team first, because you cannot architect what you have not used, and you cannot write guardrails you do not understand. Standout GCs will see and seize the chance to make AI efficient, effective, and safe across the company. The ones who lead will treat the absence of a settled rulebook not as a reason to wait, but as the opening to build the framework themselves, govern in real time, and carry the business forward through what comes next.

[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.

[Alice Davidson](#)



Securities and M&A lawyer

Tipalti

Alice Davidson is a Securities and M&A lawyer with over 15 years of experience, including deep expertise in fintech, paytech, and crypto. Before Tipalti, Davidson was Chief Legal Officer at Mogo Inc., a leading Canadian fintech, where she advised on strategic initiatives, including acquisitions, investments, financings, and partnerships, and led government relations. Over seven years, she guided Mogo from a legal, regulatory, and compliance perspective, facilitating the launch of all products, expansion across Canada, and listing on the TSX and NASDAQ. Prior to Mogo, she was an associate at Stikeman Elliott LLP, focusing on public M&A, corporate finance, and corporate governance. Davidson holds a BBA from the Schulich School of Business and a JD from the Allard

