

What GCs should consider for US AI deployment in 2026

June 2026

Executive summary

For many organizations, the question is no longer whether to adopt an AI governance program. Most organizations already have one (or at least the beginnings of one) in the form of AI use policies, intake processes, vendor diligence questionnaires, data-use restrictions, employee training and legal review procedures. But the AI landscape is moving faster than many of those programs were designed to handle.

In 2026, AI governance is becoming less about whether and how employees may use generative AI tools, and more about how organizations manage AI that is embedded across the enterprise: in software as a service (SaaS) platforms, customer-facing products, developer tools, HR systems, cybersecurity workflows, marketing stacks, productivity suites and increasingly autonomous AI agents. The result is that many organizations' original AI policies and procedures, which are often focused on employee prompting, confidential information and public chatbot use, need to be updated for a more complex environment.

This update highlights some key issues that general counsels and legal departments should be revisiting now as part of their AI governance. These include the rise of agentic AI, the expansion of third-party and SaaS vendor AI risk, topical updates regarding AI and IP (including open source and licensing of training data), AI litigation risk, and a rapidly developing patchwork of AI-specific laws and regulations in the United States and abroad.

The core takeaway is simple: AI governance should not be treated as a one-time policy project.

It should be a living legal, compliance, privacy, security and product governance framework that evolves with the technology. But AI governance can build on existing policies, rather than requiring an entire "net new" set of policies.

For GCs, 2026 is the year to pressure-test whether existing AI governance fits the ways in which AI is procured, deployed, embedded and used across the organization.

AI agents

1. Agentic AI has arrived. So too have new compliance challenges:

- AI agents, or agentic AI, refers to AI systems that autonomously execute multiple steps with minimal human involvement.
- They represent a shift in AI use from "answer generators" to workflow actors that can analyze contracts, triage customer service requests, update records, prepare first drafts, engage in commerce and execute similar multi-step processes – raising new questions around authority, supervision, auditability and privilege.
- Agents can use tools, access data and take actions on behalf of employees or customers. Unlike stand-alone chatbots, agents may connect to email, calendars, customer relationship management (CRM), HR information systems (HRIS), ticketing, procurement, code repositories and document systems, creating legal risk around access controls, scope of authority, cybersecurity, vendor management, consent, data minimization, and misuse of confidential information.
- Agentic systems create a new governance challenge: delegation without clear accountability. As AI agents become capable of making recommendations, initiating communications, negotiating terms, escalating issues or triggering transactions, GCs will need policies defining permissible uses, human-in-the-loop requirements, logging, testing, and responsibility for errors or unauthorized actions.

Major insurers are moving to exclude losses caused by AI agents from coverage. Organizations should consider consulting with their insurance brokers to understand the extent of their coverage with respect to AI agents.



AI vendors and third-party risk

2. AI vendor contracting requires new intake, assessments and contract templates:

- **Treat AI as a standard vendor-risk issue, not a special-case add-on** – Most SaaS vendors are embedding AI into existing products, copilots, analytics, support tools, personalization features and workflow automation, so vendor intake should assume AI may be present unless affirmatively ruled out.
- **Update diligence questions to capture both obvious and hidden AI uses** – Intake forms should ask whether the vendor uses AI to process customer data, train or fine-tune models, generate outputs, make recommendations, support automated decisions or enable customer-facing agentic functionality.
- **Refresh contract templates for AI-specific rights and restrictions** – Contracts for AI systems should address training use, fine-tuning, model improvement, ownership of outputs and derivatives, confidentiality, data retention, security controls, audit rights, explainability, human review, prohibited uses, incident notice and liability for AI-enabled errors. In particular, organizations also should attempt to negotiate indemnity and limitation of liability provisions that are commensurate with the risk presented by the subject AI system. An AI circuit breaker that triggers a pause or override when an AI system acts unpredictably, exceeds acceptable risk levels, or falls below a minimum performance threshold also is worth considering.
- **Do not rely on legacy data processing agreement (DPA)/ security-review processes alone** – Traditional privacy and security reviews may miss AI-specific risks, including model memorization, prompt injection, data leakage through outputs, synthetic data claims, hallucinated content, biased recommendations and downstream use of customer data for model development.
- **Build a repeatable AI vendor classification framework** – GCs should distinguish between vendors that merely use AI internally, vendors that process company data through AI features, vendors that provide customer-facing AI functionality, and vendors whose AI may influence legally or commercially significant decisions.

3. Pricing for SaaS and AI vendors is changing from seat-based to a usage- and/or outcome-based model (or a hybrid model). Legal and procurement teams need to be aware of these trends, address them in their templates and playbooks, and put processes in place to monitor pricing and policy changes:

- **Usage-based pricing** – This one was inevitable given the adoption of AI by SaaS providers and the proliferation of products that are effectively AI “wrappers.”
- **Outcome-based pricing** – Providers charging based on how many tasks are successfully completed by agents. What is a task? What is success and who defines/decides it? These and other issues need to be understood by business stakeholders and legal, and well defined in the agreement.
- **Hybrid** – Some providers are charging based on a hybrid approach of seat-, usage- and outcome-based.
- These details are often found on a dashboard that legal and procurement teams will never see or access, or in the fine print of online documents that have been updated automatically, long after a master service agreement (MSA) has been negotiated. Some organizations set up alerts for when providers change terms, policies, etc., to be aware of public-facing changes.

4. Certain SaaS and AI vendors are more prone to disruption and obsolescence than ever before. This needs to be incorporated into your third-party risk calculus.

Performance of general-purpose models is improving by multiple factors on an annual basis, while the cost of running them is decreasing. There are numerous consequences of this, including that general purpose models can now compete with, and are already competing with, specialized SaaS and AI applications. This is especially true for certain wrapper applications that are built on top of frontier models like those offered by OpenAI or Anthropic. Meanwhile, coding agents offered by Anthropic and OpenAI, and others like Cursor and Replit, provide development teams the ability to build home-grown applications and potentially obviate the need for some third-party applications. Legal teams should spend more time with their business stakeholders understanding the positioning of vendors' and their risks of disruption and, where appropriate, hedge against those risks contractually. For example, with vendors that are most prone to disruption, avoid long term commitments, exclusivity, and any other provisions that limit flexibility.

AI and IP

5. As data scraping copyright claims against foundation models continue moving through the courts, an increasing number of content creators have chosen to license their content. An established market for the data may offer content owners a stronger basis for damages in a lawsuit and dilutes a potential fair use defense for data scrapers.

A foundation model is an AI model that was trained for a wide range of tasks using a broad data set. Many foundation models were trained on data obtained through data scraping. Data scraping (or just scraping) refers to an automated process for extracting large amounts of data from websites, social media and similar online sources (and saving and using the extracted data for various purposes).

When the first foundation models hit the market, content creators quickly became aware that their copyright material was used to train these models. In one early case, output generated by a foundation model still showed the content owner's watermark.

After filing copyright infringement lawsuits, many large content creators recognized the business opportunity of licensing their content as training data as a practical alternative. As a consequence of these licensing opportunities, a business acquiring data through scraping faces a potentially greater copyright infringement risk, because the commercial availability of data dilutes a potential fair use defense (specifically, the market effect prong).

In addition to copyright infringement, data scraping also presents other legal risks: breach of contract (violating the terms applicable to the source of the scraped data), technical controls on the scraped website (e.g., whether a robots.txt file is on the scraped website), whether personal information is part of the scraped data, how and when scraping occurs, and how the scraped data is used.

Major developers offer a range of indemnities to their commercial licensees. Development of application built on underlying models may be less resonated to do so with respect to what they have added. Further, your organization itself may be looking to train an AI application on third-party data. To help manage the risks of doing so, consider some guardrails, such as:

- Do not conduct data scraping while logged into a website and do not circumvent technological measures that limit or prevent access to certain areas of a website. Generally, use of "publicly available" data – including publicly available personal data – is lower risk. Good recordkeeping is key: the business must be able demonstrate that the scraped data was publicly available when scraped and not subject to contractual or other restrictions, and that scraped personal data met the definition of publicly available under privacy laws or was deidentified before use. Even with these guardrails in place, scraping publicly available data is not risk free because publicly available data still is subject to copyright laws.

- Avoid scraping or otherwise acquiring data that is potentially competitively sensitive, particularly if nonpublic. Competitive intelligence data can trigger claims under consumer protection and competition laws, particularly when shared among competitors and used for pricing or similar purposes that can harm consumers. For instance, in 2025, algorithmic data-sharing platform provider Realpage settled with the Department of Justice regarding allegations that its service enabled competitors to share and combine nonpublic, competitively sensitive apartment rental data to collude on suppressing supply and aligning pricing. Both regulatory and consumer claims were also brought against the platform users (i.e., the landlords).
- Conduct data scraping when unlikely to interfere with the website's business (e.g., early morning or late-night local time). Delay accessing multiple pages on the same domain. Add idle time between requests. Limit the depth of crawl within the given domain. Minimize volume of requests to a domain. Only scrape the parts of pages required for the purpose.

Also, keep in mind that the US Copyright Office has determined that, in most cases, AI outputs will lack sufficient human expression to qualify for copyright protection.

AI regulatory horizon

6. AI-specific laws in the US: the US federal government is in a tug of war with states regarding the (de)regulation of AI. The proliferation of general, omnibus AI legislation seems unlikely in view of this dynamic; that said, there are a number of narrower AI-specific laws in effect in the US.

- The Trump Administration issued an executive order entitled "[Removing Barriers to American Leadership in Artificial Intelligence](#)" on January 23, 2025, three days after inauguration. As its title indicates, the Executive Order explicitly frames AI development as a matter of national competitiveness and economic strength by prioritizing removal of what the Administration views as regulatory obstacles to innovation. Since then, the Administration has actively promoted its agenda with an AI Action Plan, a US\$90 billion investment to accelerate the development of an AI infrastructure in Pennsylvania (July 15, 2025) and several other [Executive Orders](#). In March 2026, the White House released its [National Policy Framework for Artificial Intelligence](#) to guide Congress. This framework calls on Congress to create a unified national approach to AI and "preempt state AI laws that impose undue burdens," but specifically carve out "state laws of general applicability ... particularly laws to protect children, prevent fraud, and protect consumers," as well as zoning laws for data centers and requirements for a state's own use of AI.

The Administration also has threatened states implementing AI laws with cuts to their federal government funding. And in April 2026, the Department of Justice recently intervened in a case in which the Colorado AI Act was challenged on constitutional grounds ([xAI v. Weiser](#) (D. Colo., 1:26-cv-01515)).

The Administration appears to be reconsidering this approach, at least with respect to advanced AI models in light of concerns regarding the potential impacts on cybersecurity from misuse of Anthropic's "Mythos" product (discussed further below) and on June 2, 2026, President Trump signed an executive order [[Promoting Advanced Artificial Intelligence Innovation and Security – The White House](#)] providing federal government hardening of cybersecurity defenses against AI and prioritizing enforcing existing cybercrimes laws, as well as directing a multiagency effort to develop a voluntary program for covered frontier models to be assessed prerelease for cyber risk to critical infrastructure. This light touch security-focused approach is consistent with the federal government's ongoing prioritization of innovation over regulation, but with a new focus on critical security risk.

Meanwhile, states continue to pass more AI-related laws that regulate "significant decisions" by AI (e.g., the California Consumer Privacy Act (CCPA)) and the use of automated decision-making technology systems and/or AI in the HR context, and AI use in other sectors or industries, including education, healthcare and real estate. Other state laws relate to transparency and safety, for LLM providers with significant numbers of users, chatbot companions and deepfakes. Of course, these are only representative of some examples, and types of AI-specific laws and regulations in the US. States have passed dozens of AI-related bills, creating a patchwork that rivals the state consumer privacy laws. The International Association of Privacy Professionals (IAPP) maintains a [state AI legislation tracker](#) (as well as a global AI legislation tracker). In mid-May, bipartisan House efforts have emerged to specifically preempt state AI safety bills, at least for two years, though the scope of federal regulation to be included in such a bill remains subject to debate. Prior efforts at comprehensive federal AI legislation have failed.


7. Cybersecurity and incident response remain a key consideration, particularly for agentic AI.

Cybersecurity is a key consideration for agentic AI because agents not only generate outputs, but also access systems, retrieve sensitive data, call tools, send communications, execute transactions or modify records – expanding the potential attack surface and consequences of compromise. Organizations must evaluate agentic AI through a cyber-risk lens, including access controls, least-privilege permissions, prompt-injection and data-exfiltration risks, logging, monitoring, human approval thresholds, incident response, and vendor/security diligence.

Anthropic's [Claude Mythos Preview](#) illustrates why cyber-risk is central to agentic AI governance: Anthropic reported that its new model, Mythos, can identify and exploit zero-day vulnerabilities in major operating systems and browsers, and that nonexperts were able to use it to generate working remote-code-execution exploits. While Mythos' rollout is proceeding more prudently, Anthropic has said it does not plan to make Mythos Preview generally available. This said, its reported capabilities, which other models will likely possess in the near future, highlight a broader GC-level concern: frontier agents may compress vulnerability discovery and exploitation timelines, weakening security controls that rely on attacker friction rather than hard barriers.

8. Consider data subject rights and other consumer protection matters. Privacy laws in the US, EU/UK, and rest of world have historically been, and will continue to be, one of the main drivers of AI regulatory enforcement. AI is also subject to general purpose consumer protection and safety laws, as are other products and services:

- If an AI system will process personal data, consider both the basis for the use of the personal data, as well as how data subject rights requests for access, objection to processing, and deletion/erasure can be honored.
- State consumer protection, labor and/or privacy laws may also require certain notices, and, in some circumstances, require certain opt-out, appeal and/or information rights regarding AI that uses personal data to make (or in some cases assist in making) significant decisions affecting an individual.
- US privacy regulators may focus on how data minimization principles apply to use of personal data to train AI models. Notably, some prominent AI vendors have released tools to filter personal data from prompts before it is sent to AI models, which are designed to help organizations comply with privacy laws. Organizations should do diligence on these sorts of third-party tools, and also implement processes and procedures internally to minimize the amount of personal data ingested into AI tools for training and other purposes that may be considered by regulators to be secondary to providing consumers with products and services. To the extent personal data is needed for improving products or services, including for algorithmic training, look to disclose those purposes at collection.
- The US plaintiffs' bar is already coming after AI, including under confidentiality of communications laws like the [California Invasion of Privacy Act](#) (CIPA), and under consumer protection and product liability laws. Make sure that AI use cases, especially those that are consumer-facing, undergo product counseling review.
- The Attorney Generals in California, Massachusetts, New Jersey, Oregon and Texas have all issued specific guidance on the application of their state consumer protection and civil rights laws to AI claims and use cases. State and federal authorities can, and can be expected to, bring claims under existing laws concerning deception, unfairness or discrimination arising out of AI promotion and use.



Contacts

For more information on AI and how to develop, implement or refine an ethical AI policy and framework for your business, contact:

US



Alan Friel
alan.friel@squirepb.com



Julia Jacobson
julia.jacobson@squirepb.com



Kyle Fath
kyle.fath@squirepb.com



Glenn Brown
glenn.brown@squirepb.com



Kyle Dull
kyle.dull@squirepb.com

Asia Pacific

China



Lindsay Zhu
lindsay.zhu@squirepb.com

Hong Kong



Nick Chan
nick.chan@squirepb.com

Japan



Scott Warren
scott.warren@squirepb.com

EMEA

UK



David Naylor
david.naylor@squirepb.com

Germany



Dr. Annette Demmel
annette.demmel@squirepb.com

Spain



Bartolomé Martín
bartolome.martin@squirepb.com

MiddleEast



Habib Saeed
habib.saeed@squirepb.com

	Powered by Squire Patton Boggs
	Privacy World Blog
	AI Law & Policy Hub

The opinions expressed in this update are those of the author(s) and do not necessarily reflect the views of the firm, its clients, or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

© Squire Patton Boggs. All Rights Reserved 2026.