AI Is Already in Your Business. Is Your Legal Team Ready?

By Samir A. Bhavsar, Partner, Baker Botts

Are you confident you know every AI tool your company is using right now, including the ones your employees adopted on their own? Do you know whether your AI vendor contracts actually protect your company's data, or whether they quietly grant the vendor the right to train their models on your proprietary information? If one of your company's AI-driven decisions triggered a discrimination claim tomorrow, could your team explain how that decision was made? If any of these questions give you pause, you are not alone, and you are exactly who this article is written for.

I recently had the pleasure of presenting a CLE to ACC-DFW on proactive AI risk management, and the questions I received from in-house counsel during and after that session made clear that there is strong demand for practical, actionable guidance on this topic. To continue that conversation, **I am offering complimentary, structured 60-minute AI governance briefing sessions** to ACC-DFW member companies and their legal departments. These sessions are designed to give your team a candid, confidential assessment of where your AI risk exposure may lie and what a practical compliance roadmap may look like for your organization. Details on how to schedule a session are at the end of this article.

The Regulatory Landscape Is Not Waiting for You

Regardless of your industry or the size of your legal department, the compliance clock is already running. The global regulatory environment for AI has shifted from aspirational guidance to enforceable obligations with real financial consequences, and the timeline is tighter than most in-house counsel appreciate.

The most significant development on the global stage is the EU AI Act, the world's first comprehensive AI regulatory framework. It applies not just to European companies, but to any company whose AI systems are placed on the EU market or affect individuals in the EU. If your company has European customers, employees, or business partners, this law almost certainly applies to you. Penalties for violations involving prohibited AI practices reach up to €35 million or 7% of worldwide annual revenue, and separate penalties for high-risk system violations can reach €15 million or 3%. The compliance deadline for high-risk AI systems is August 2, 2026 — not far away — and building a compliant AI governance program typically takes a few months, which means companies that have not started are already behind.

The domestic picture is equally demanding, though structured differently. Rather than a single federal framework, the United States regulates AI through an expanding patchwork of sectoral laws and state-level legislation. The FTC is actively using its consumer protection authority to pursue deceptive and discriminatory AI practices, as illustrated by its ban on the use of biased facial recognition technology in retail locations. The EEOC is applying existing employment law

to AI-driven decisions that affect job applicants and employees. And the states are moving fast: Texas's Responsible AI Governance Act took effect January 1, 2026, Colorado's AI Act takes effect June 30, 2026, and California continues to expand its AI-specific regulatory requirements. If your company operates across multiple states, or globally, you are not dealing with one compliance obligation. You are dealing with many overlapping ones, and a one-size-fits-all approach will not work.

It is also worth noting that not all AI carries the same level of regulatory exposure. Most of the major frameworks described above are primarily concerned with AI systems that make or influence consequential decisions affecting people, such as your customers and employees — what practitioners call external-facing AI. Internal productivity AI tools like carry real risk too, but a different kind of risk.

The key takeaway for in-house counsel is straightforward: the regulatory regime that applies to your company depends on your industry, your AI use cases, and your geographic reach. Understanding that intersection is the starting point of every meaningful AI governance conversation.

This Is a Board-Level Issue

One of the things that seemed to resonate most during the CLE is that AI governance is no longer an IT issue or an innovation initiative. It is a core governance risk, and Boards of Directors are increasingly expected to treat it as such.

AI systems today influence revenue, pricing, employment decisions, customer access, and safety. When those systems fail, through bias, hallucination, or data leakage, the consequences can include regulatory enforcement, class action litigation, reputational damage, and operational disruption. Regulators, investors, and insurers now expect Board-level oversight of AI risks, just as they expect oversight of financial controls and cybersecurity. Courts are beginning to ask whether Boards exercised reasonable oversight over AI systems when failures occur, and the absence of that oversight is increasingly being framed as a governance failure, not a technical one.

In-house counsel play a critical role here. Your job is to translate AI activity into legal and litigation risk the Board can act on, ensuring the Board knows where AI is material, that management has appropriate controls in place, and that someone is accountable if something goes wrong. In my practice, the question I hear often from in-house teams is not “do we have AI risk?”, it is “where do we even start?” You are the internal escalation point, and the documentation you create around Board-level AI oversight is itself a meaningful form of protection under the business judgment rule.

AI Is Creeping Into Your Business in Ways You May Not Realize

Here is something I find comes as a genuine surprise to many in-house lawyers: the most significant AI risks often come from sources that no one in the legal department approved, reviewed, or even knows about.

Consider a few scenarios drawn directly from recent litigation. A fitness equipment company integrated a third-party AI chatbot on its website for customer service, and that chatbot captured and analyzed customer conversations without explicit consent, generating California wiretap claims. A major health insurer deployed an AI system that automated nearly 300,000 insurance claim denials in two months, triggering coordinated class action lawsuits alleging ERISA violations and bad faith claims handling. A fast-food chain rolled out AI-powered facial and voice recognition at drive-throughs without proper biometric consent procedures or demographic bias testing, resulting in BIPA claims in Illinois. A mental health chatbot deployed without adequate safety controls engaged with a vulnerable teen user without any crisis escalation protocol in place, and the company now faces a wrongful death lawsuit.

None of these companies set out to create these problems. In each case, someone adopted an AI tool because it seemed useful, without considering the legal implications until after the harm occurred. That is the pattern proactive AI governance is designed to prevent.

It helps to draw a distinction that will shape the rest of this article. AI tools in a business context fall into one of two broad categories. Internal-facing AI consists of tools that help your employees work, such as drafting assistants, document summarization platforms, internal knowledge search tools, and general-purpose AI tools that employees use to speed up everyday tasks. The risks are real but primarily operational and confidentiality-based: employees inadvertently sharing sensitive data with public AI platforms, proprietary information being used to train models competitors can also access, and trade secrets leaking in ways that are difficult to detect.

Enterprise "shadow AI," meaning AI tools that employees adopt on their own without official approval or oversight, grew nearly 5x in a single year. Over half of employees using AI on the job report receiving no training whatsoever on AI risks. And more than one-quarter of what employees paste into public AI tools turns out to be sensitive corporate data.

External-facing AI, by contrast, is where regulators, enforcement agencies, and plaintiff's lawyers are directing most of their attention. These are AI systems embedded in your products, services, and decision-making processes that directly affect customers, employees, job applicants, and the general public. When an AI system decides whether to approve or deny an insurance claim, screens a job applicant in or out of consideration, prices a product for a particular customer, or responds to a consumer inquiry, it is making or influencing a consequential decision about a real person's access to employment, credit, insurance, healthcare, or services. That is the category the EU AI Act classifies as high-risk. That is what the FTC, the EEOC, and state regulators are actively pursuing. And that is where class action plaintiffs are seeking multi-million-dollar verdicts . Every one of the case studies described above involves external-facing AI. That is not a coincidence.

The Common AI Risk Hotspots

While a full breakdown of every AI risk area is beyond the scope of this article, a few categories appear most consistently across industries. Some fall into the internal-facing category, others into the external-facing category. As noted above, it is the external-facing category that draws the most intense regulatory and litigation scrutiny. Which categories are most urgent depends heavily on industry, geographic footprint, and the specific ways AI has entered your operations. For example,

a financial services company, a healthcare provider, and a manufacturer may face nominally the same categories, but their priorities, regulatory frameworks, and remediation needs will look very different.

Vendor contracts are one of the most urgent. Most standard AI vendor agreements are written to favor the vendor, including provisions that allow the vendor to use your company's data to train their AI models, vague "aggregate data" carve-outs that sound harmless but are not, and liability caps that leave your company exposed when the vendor's AI causes harm. Negotiating these contracts requires specific AI governance expertise, and most form agreements will not get you where you need to be without significant revision.

Employment and HR applications carry particularly high legal exposure. AI tools used in hiring, resume screening, workforce management, and HR decision-making are subject to federal anti-discrimination law under Title VII, the ADA, and the ADEA, as well as a growing body of state and local rules. Courts have already held that employers cannot outsource liability for discriminatory AI hiring tools to the vendor; if the AI screens out candidates in a discriminatory way, the employer may be responsible. That remains true even when working with a reputable vendor, because vendor bias testing is rarely performed on the employer's own data.

Consumer-facing AI such as chatbots, recommendation engines, pricing algorithms, and content personalization tools, creates disclosure obligations, consumer protection exposure, and product liability risk that many companies are not yet accounting for. Data privacy intersects with virtually every customer-facing AI use case, and the overlap between AI governance and existing privacy law, particularly GDPR and California's CPRA, creates a dual compliance challenge requiring integrated analysis.

The point is not that every company faces these risks equally. You cannot know which are most pressing without first understanding how AI is being used across your business, which is exactly what the briefing sessions at the end of this article are designed to help you assess.

A Practical Path Forward: Phase-Based AI Governance

For companies that are just beginning to build an AI governance program, or that have AI activity happening across the business without a structured framework, I recommend a three-phase approach.

Phase 1 is discovery. Before you can manage AI risk, you have to know where AI is being used. That means auditing every department for AI tools in use, including both officially approved tools and the ones employees are using on their own. In my experience, this audit almost always surfaces AI tools the legal department had no idea were in use, and that realization tends to be the moment when AI governance shifts from an abstract compliance obligation to a real operational priority. The output is a complete, current picture of your AI footprint, including a data flow map, a compliance matrix oriented to your specific geographies and sectors, and a reviewed inventory of your vendor contracts.

Phase 2 is risk profiling and compliance. With that picture in hand, you assess risk, conduct privileged bias audits on AI systems already in use, and build the documented compliance

framework — including, where appropriate, documentation aligned with the NIST AI Risk Management Framework — that regulators and courts will expect to see if something goes wrong.

Phase 3 is continuous assurance. AI governance is not a one-time project. It requires ongoing monitoring, regular bias audits, and a program built to adapt as new regulatory demands emerge. The companies that invest in building durable governance infrastructure now will be the ones best positioned to demonstrate to regulators, insurers, and courts that they took their obligations seriously.

The cost of getting this wrong is significant. Regulatory fines under the EU AI Act alone can reach €15 million or 3% of global revenue for high-risk AI use cases. Class action settlements in AI cases can reach into the millions. The operational disruption that comes from a forced AI system shutdown or a government investigation is difficult to quantify but impossible to ignore. Put simply, the cost of non-compliance can easily run 50 to 100 times the cost of building a proper governance program in the first place.

Schedule Your Complimentary AI Governance Briefing

If your company is using AI, or if you are not sure whether it is, now is the right time to take stock. Building a governance program does not have to be complicated, and you do not have to figure it out alone. I am offering **complimentary 60-minute AI governance briefing sessions** specifically for ACC-DFW member companies and their legal departments. Each session is a structured conversation tailored to your situation, covering the regulatory landscape relevant to your industry, the most common AI risk hotspots for companies like yours, and what a realistic governance roadmap looks like for your current AI footprint. These are not sales pitches. They are practical working sessions, and you will leave with something useful. I have a limited number of sessions available each month and am prioritizing ACC-DFW members on a first-come, first-served basis.

To schedule your session, please contact me directly:

Samir Bhavsar

214.953.6581

samir.bhavsar@bakerbotts.com

Samir Bhavsar is a partner in the Dallas office of Baker Botts, where he is a member of the Privacy, Cybersecurity, and AI Governance practice. Samir is one of a small number of lawyers in the United States who holds all three of the International Association of Privacy Professionals' most rigorous practitioner credentials: the Artificial Intelligence Governance Professional (AIGP), the Certified Information Privacy Professional for the United States (CIPP/US), and for Europe (CIPP/E). That combination — AI governance framework expertise, US domestic regulatory knowledge, and European data protection depth — positions him to advise companies navigating the converging demands of AI regulation across jurisdictions.