Legal Lines Around AI

A 6-Part Series

MAC MURRAY & SHUSTER

PART 1

Navigating the AI Legal Landscape

Welcome to **Legal Lines Around AI**, a six-part blog series exploring how AI laws are taking shape in the United States and what those changes mean for businesses using, building, or relying on AI systems. Throughout this series, we'll break down emerging legal requirements, highlight risk triggers, and offer practical guidance to help organizations navigate an increasingly complex and fragmented AI governance environment.

While AI regulation comes in many forms, this series focuses primarily on comprehensive, cross-sector AI laws—not industry-specific rules governing areas like health care, employment, or AI companion chatbots. Our goal is to help businesses understand the broad frameworks that are most likely to apply across operations, technologies, and use cases.

Before diving into specific obligations, it's important to understand the legal backdrop shaping AI governance today.

A Shifting Federal Approach to AI

At the federal level, AI policy has undergone a dramatic shift in just a few years.

In October 2023, the Biden Administration issued Executive Order 14110, the most sweeping federal AI action to date. That order directed federal agencies to address AI-related risks tied to safety, civil rights, consumer protection, privacy, and the government's own use of AI systems.

That approach changed course in January 2025, when the Trump Administration rescinded Executive Order 14110 and adopted a more innovation-first, deregulatory posture, emphasizing U.S. competitiveness in AI development. A follow-up executive order issued in December 2025 reinforced that shift, signaling a preference for minimal regulatory burden and raising the possibility of future federal preemption of state AI laws, but stopping short of establishing a comprehensive federal AI statute.

Regulation Without a Federal AI Law

Even without a dedicated federal AI law, AI regulation hasn't stopped.

Federal agencies have continued to apply existing authority to AI-related practices. The Federal Trade Commission, for example, has made clear that it will use its unfair and deceptive acts and practices authority to police AI claims, data practices, bias, and consumer harm. In practice, that means AI systems are already subject to scrutiny, even in the absence of new legislation.

States Step In—and Move Fast

With no comprehensive federal framework in place, states have become the primary drivers of AI regulation.

Early state regulation often came through comprehensive privacy laws, which introduced direct oversight of profiling, generally defined as the automated processing of personal data to evaluate or predict aspects of an identified or identifiable consumer, such as economic circumstances, health, preferences, behavior, reliability, location, or movements.

These laws impose heightened obligations when profiling is used to produce "legal or similarly significant effects," including decisions affecting access to financial services, employment, housing, insurance, education, health care, criminal justice outcomes, or other essential goods and services. Those obligations may include consent requirements, opt-out rights, and mandatory risk assessments. Colorado and California went a step further by adopting comprehensive automated decision-making regulations under their respective comprehensive privacy laws, adding more structure and specificity to these requirements.

Beginning in 2023, and accelerating through 2025, several states expanded beyond privacy laws, enacting both sector-specific and cross-sector AI legislation addressing automated decision-making, algorithmic discrimination, deepfakes, consumer interactions, employment screening, children's data, and broader consumer protection concerns. The result is a rapidly evolving and increasingly complex patchwork of state AI requirements.

Why "High-Risk" AI Matters

Many of the most significant AI obligations turn on how an AI system is used and who is responsible for it.

States like Colorado and California distinguish between organizations that develop or substantially modify AI systems (such as a technology company that designs and trains a résumé-screening algorithmic system) and those that deploy AI systems (such as an employer that uses that tool to evaluate job applicants). In many cases, obligations are triggered only when an AI system is deemed "high risk."

Colorado's most comprehensive AI law applies broadly across entities, while California's most comprehensive AI requirements apply to organizations that meet the definition of a "business" under the CCPA. Understanding where your organization fits – and whether your AI systems fall into a high-risk category – is often the key to determining which legal requirements apply.

What Businesses Should Be Doing Now

Even as the legal landscape continues to shift, one takeaway is already clear: AI governance starts with visibility.

Businesses using AI should:

- Inventory where and how AI systems are used across the organization, and
- Assess whether those systems make, or materially influence, significant decisions about consumers, employees, or applicants.

That threshold often determines whether heightened legal obligations apply.

PART 2

Are You Developing or Deploying a "High-Risk" AI System?

Under comprehensive AI laws, one question largely determines a company's regulatory exposure: are you developing or deploying an AI system that qualifies as "high risk"?

In this second installment of **Legal Lines Around AI,** we take a closer look at how emerging AI laws draw that line, why the high-risk designation matters, and how it can fundamentally reshape a business's legal, operational, and governance approach to AI.

Developers vs. Deployers: Why the Distinction Matters

As discussed in the first post in this series, comprehensive AI laws draw a critical distinction between developers and deployers of AI systems.

- A developer is the entity that builds an AI system or intentionally and substantially modifies an existing one.
- A deployer is the entity that puts the AI system into use.

Not every system change turns a deployer into a developer. A substantial modification typically means a material change that affects an AI system's outcomes, risks, or decision-making impact. Routine maintenance, updates, or bug fixes generally do not rise to that level. This distinction matters because legal obligations and potential liability often differ significantly depending on an organization's role.

What Is a "Significant Decision"?

The concept of high-risk AI is grounded in the real-world impact of an AI system's decision-making, not the sophistication of the technology itself.

Rather than regulating all AI systems uniformly, comprehensive AI laws focus heightened requirements on systems that make or play a substantial role in making significant decisions affecting individuals' access to, or the cost or terms of:

- Financial and lending services
- Employment and workplace opportunities
- Housing
- Insurance
- Education
- Health care
- Legal services or criminal justice outcomes
- Essential goods and services.

This risk-based approach enables lawmakers to focus regulatory scrutiny on the AI systems that may be more likely to result in discrimination, unfair treatment, or privacy harm while preserving flexibility for lower-risk uses and avoiding unnecessary constraints on AI innovation.

Defining "High-Risk" AI Systems

Most comprehensive AI laws define an AI system broadly as a machine-based system designed to achieve explicit or implicit objectives by inferring from inputs how to generate outputs, such as predictions, recommendations, or decisions. An AI system becomes "high risk" when it makes, or is a substantial factor in making, significant decisions.

Importantly, an AI system does not lose its high-risk status simply because a human remains "in the loop." If an AI system's outputs meaningfully influence a significant decision, heightened legal obligations may still apply, even where a human is involved. That said, certain requirements may be reduced or modified when there is meaningful human involvement in the decision-making process. To qualify as meaningful, human involvement must involve the exercise of independent judgment, not mere rubber-stamping of AI outputs. In practice, this typically requires training decision-makers to understand and critically assess AI recommendations, along with clearly documented authority to override, modify, or reject AI-driven outcomes.

Classification Is the Critical First Step

From a risk-management perspective, accurate classification of the AI system is essential.

Businesses should evaluate whether any AI systems used across the organization could be considered high risk. That analysis should include:

- The nature of the decisions involved
- Whether AI outputs materially influence outcomes
- Who controls system design, training, and modification

Misclassifying a high-risk AI system as low risk may create regulatory exposure down the line, particularly as enforcement activity increases.

What Happens Once an AI System Is "High Risk"?

Once an AI system is classified as high risk, legal and operational expectations change.

Organizations should implement robust oversight measures, including:

- Documenting the system's decision-making logic
- Regularly testing for accuracy, bias, and unintended impacts
- Clearly defining when and how human review applies
- Training employees on how to interpret and appropriately rely on AI outputs, including when human judgment should override the system

How Obligations Scale with Risk and Role

In addition to internal governance, applicable laws impose specific compliance obligations depending on whether an organization is acting as a developer or deployer, and whether the system is high risk. These obligations general scale as follows:

Role	Website Disclosures	Pre-Interaction Disclosure	Adverse Post-Use Notice	Consumer Rights	Risk Assessments	Prohibited Behaviors
Developer of AI Systems	Yes	Yes	No	No	No	Yes
Deployer of AI Systems	No	Yes	No	No	No	Yes
Developer of High-Risk AI Systems	Yes	Yes	No	No	No	Yes
Deployer of High-Risk AI Systems	Yes	Yes	Yes	Yes	Yes	Yes

For deployers of high-risk AI systems, obligations often expand to include consumer rights, post-decision notice requirements, and formal risk assessments. We will dive deeper into these requirements throughout the **Legal Lines Around AI** series.

PART 3

The Expanding Disclosure Obligations for AI Systems

Disclosures are a central regulatory requirement across consumer protection laws to promote transparency, fairness, and accountability.

In this third installment of **Legal Lines Around AI**, we examine how emerging AI laws are expanding disclosure requirements, why those obligations matter from a risk perspective, and how businesses can build AI governance programs that keep pace with evolving transparency expectations.

AI Disclosures Build on Familiar Consumer Protection Principles

AI disclosure obligations may feel new, but the underlying principles are not.

Comprehensive privacy laws, for example, require organizations to provide clear, upfront notices, including disclosures related to profiling, that generally explain what data is being collected, how it will be used, and what, if any, rights may be available to individuals. Telemarketing laws take transparency a step further by imposing real-time disclosure requirements, such as identifying the caller and the nature of the call at the start of a conversation and again before completing a transaction.

Emerging AI laws follow a similar path, with some states applying these transparency principles across the lifecycle of an AI-driven interaction.

Multi-Stage Disclosure Across the AI Lifecycle

Rather than relying on a single notice, some states are implementing multi-stage disclosure regimes that trigger obligations at defined checkpoints across the AI lifecycle and consumer interaction.

While the specific content and triggers vary by jurisdiction, the shared goal is consistent: to ensure individuals are not unknowingly subject to AI-driven decision-making that affects their rights, opportunities, or access to essential services, and to clearly signal when automation is shaping an interaction.

Common disclosure checkpoints include:

- **Public website disclosures** describing AI use at a high level
- **Privacy policy disclosures** explaining how personal information will be used in AI systems
- **Pre-use notices** provided before an AI system is applied to an individual
- **Real-time interaction disclosures** when consumers directly engage with AI
- **Post-use notices**, in some jurisdictions, explaining how AI influenced a decision or outcome

Public Website Disclosures

Public-facing disclosure obligations typically fall on developers of AI systems; that is, entities that build or intentionally and substantially modify AI systems.

Under California AI law, developers must publish a high-level summary describing the training data used to develop an AI system. While the law stops short of requiring disclosure of raw datasets or sensitive technical details, it does mandate transparency regarding the nature, source, composition, and treatment of training data. This includes whether the data contains personal information, protected intellectual property, or synthetic elements; how and when it was collected and used; and how it supports the AI system's intended purpose.

Colorado's AI law goes further for high-risk AI systems, requiring developers to disclose via a public use-case inventory or their website the types of systems they build or substantially modify and how they identify and manage reasonably foreseeable risks of algorithmic discrimination. Developers must also provide deployers with detailed documentation covering intended and potentially harmful uses, training data summaries, known limitations, discrimination risks, evaluation methods and data governance practices, mitigation measures, and guidance for proper use.

Importantly, Colorado law also requires developers to report known or reasonably foreseeable risks of algorithmic discrimination to the Colorado Attorney General and to all known deployers or developers within 90 days of discovery.

Privacy Policy Disclosures

Comprehensive privacy laws increasingly require organizations meeting applicable thresholds to explain, in their privacy policies, when and how they use profiling that may produce legal or similarly significant effects on individuals.

Colorado's privacy law sets the most onerous standards, requiring clear disclosures about:

- The logic behind such automated decision-making
- The types of data used
- The purpose of the profiling
- The rights individuals have to understand, challenge, or opt out of decisions that meaningfully affect access to services, opportunities, or benefits

These requirements create pressure to align legal disclosures closely with how AI systems function in practice.

Real-Time Interaction Disclosures

Several states now require businesses to notify individuals when they are interacting with AI—but each takes a different approach to when and how that disclosure must occur.

California, Colorado, Maine, and Utah generally require disclosure at the start of an interaction when AI is used for consumer engagement, but with notable variations:

- **California** treats disclosure as a safe harbor to its prohibition on deceptively presenting AI as human.
- **Colorado** waives disclosure when the AI nature of the interaction would be obvious to a reasonable person.
- **Utah** generally requires disclosure for most businesses only if a consumer asks, but individuals providing services in regulated occupations must prominently disclose when a person is interacting with AI in a high-risk interaction.
- **Maine** requires disclosure only when necessary to avoid misleading a reasonable consumer into believing they are interacting with a human.

California's telemarketing law adds another layer, requiring that, before an AI-generated message is delivered, a live caller must state the nature of the call and the identity of the

business, obtain the recipient's consent to hear the prerecorded message, and clearly disclose that the message uses an artificial voice.

Pre-Use Notices for High-Risk AI

Both California and Colorado require pre-use notices when a high-risk AI system is used to make, or play a substantial role in making, a significant decision about an individual.

This notice must be delivered prominently at or before data collection, or before previously collected data is repurposed for use in a high-risk AI system. Required disclosures must clearly explain:

- The specific purpose of the AI use
- The nature of the decision being made
- How to find additional information on the deployer's website
- The consumer's rights to access or opt out of the AI system
- The prohibition on retaliation
- How the system processes personal information, what outputs it generates, and how decisions will be made if the consumer chooses to opt out

Post-Use Notices Following Adverse Decisions

Colorado's AI law also requires post-use notices when a high-risk AI system is used to make, or meaningfully influence, a significant decision that has an adverse outcome to the consumer.

The deployer must explain the principal reasons for the outcome, including how the AI system contributed to the decision, what types of data it processed, and where that data came from. Consumers must be given an opportunity to correct any inaccurate personal information the system relied on, along with a meaningful appeal mechanism that, when technically feasible, includes review by a human decision-maker.

Managing Disclosure-Related Risk

Outdated, inconsistent, or inaccurate disclosures can create significant compliance and enforcement risk.

To mitigate that risk, businesses should centralize ownership of AI disclosures and ensure that website statements, privacy notices, and consumer-facing communications accurately reflect how AI systems operate in practice. Legal, compliance, and technical teams should collaborate to validate disclosures with detailed system analysis and real-world disclosure/notification use cases, particularly when systems are updated or repurposed.

Treating Disclosures as Living Obligations

AI disclosures should not be treated as "set it and forget it" statements.

Establishing internal review processes tied to system changes, retraining events, or new deployment contexts can help ensure disclosures remain accurate and compliant over time. Using plain language, accessible formats, and consistent delivery methods can further reduce the risk that disclosures are misleading, insufficient, or inaccessible to affected individuals.

PART 4

The Rise of Consumer Control Over AI Decisions

In our last **Legal Lines Around AI** post, we explored how emerging AI laws increasingly rely on disclosure obligations as a front-line consumer protection tool.

But this is only the starting point. Across these same laws, disclosure requirements act as a gateway to a growing set of substantive consumer rights that attach once high-risk AI systems make significant decisions. In this post, we examine the consumer rights triggered by AI use and what those rights mean for businesses deploying AI systems in consequential decision-making.

Comprehensive Privacy Laws Anchor Consumer Rights

Comprehensive state privacy laws provide the foundational consumer rights framework that already governs personal data and many AI systems used in significant decision making.

At a high level, these laws grant individuals the right to understand how their information is collected and used; to access, correct, and delete personal data; and to opt out of the sale of personal information and targeted advertising.

More importantly for AI governance, comprehensive privacy laws anchor two rights that directly apply to high-risk AI systems:

- The right to access information about the use of a high-risk AI system; and
- The right to opt out of certain high-risk AI uses.

Together, these rights give consumers meaningful control over how their personal information is used in AI-driven decisions and place new operational demands on businesses.

Access Rights Shift from Data-Centric to Decision-Centric

Privacy and AI laws increasingly equip consumers with the right to access meaningful information about how businesses use their personal information in high-risk AI systems. However, while traditional access rights are largely data-centric (i.e., what information the business holds and where it came from), AI-specific access rights are decision-centric.

Laws in states such as California and Colorado require deployers of high-risk AI systems to provide meaningful explanations in response to a high-risk AI access request, including:

- Why a high-risk AI system was used in relation to a particular consumer
- How the system processed the consumer's data to generate an output
- How that output influenced a significant decision, including the role of any human review

Under California law, if a business plans to reuse high-risk AI outputs for additional significant decisions, it also must explain how those outputs will be used and whether human review is involved.

Minnesota and Connecticut laws further expand access rights when high-risk AI systems are used, giving consumers the right to question the outcome of an AI decision and be informed of what actions might have led to a different result, as well as what steps could be taken to influence future decisions. Consumers are also entitled to review the personal information used by the system. If the decision is determined to have been based on inaccurate personal information, consumers have the right to correct that information and to have the decision reevaluated using the updated data.

As shared previously, Colorado's AI law imposes additional obligations when a high-risk AI system produces an adverse consumer decision. In those cases, the deployer must provide a statement of the principal reason or reasons for the decision, offer the consumer an opportunity to correct inaccurate personal information, and provide a mechanism to appeal the adverse decision under human review.

These expanded access rights are significant because they require businesses to translate technical system behavior into clear, consumer-ready explanations and to maintain processes capable of changing outcomes based on new or corrected data.

Opt Out Rights as a Core Control on High-Risk AI Use

Most comprehensive state privacy laws provide an opt out right for high-risk AI systems. While many states offer this right in a relatively high-level form, California and Colorado impose the most detailed requirements and conditions.

Under California regulations, businesses generally must honor consumer opt out requests for high-risk AI uses, subject to limited exceptions. A deployer is not required to offer an opt out where:

- A clear and simple appeal process exists that includes human review with authority to overturn the AI-driven decision
- The system is used solely to assess work performance and does not result in unlawful discrimination
- The system is used solely to allocate work or determine compensation, again provided it does not unlawfully discriminate.

Colorado similarly grants consumers the right to opt out of high-risk AI systems, unless the processing involves "meaningful human involvement." This is defined as a substantive, independent review of the data or AI output by a human with authority to change or influence the resulting decision.

When consumers opt out of high-risk AI processing, deployers face specific response and process obligations. Businesses must:

- Provide confirmation that the opt out request was honored
- Allow opt-outs by specific use case, provided a single, universal opt out option is also available covering all uses

Timing matters. If an opt out request is submitted before processing begins, the AI system must not be used. If the request comes after processing has started, the deployer must stop AI processing within 15 business days.

Requests may be denied if reasonably believed to be fraudulent, but the basis for denial must be explained. Even re-consent is also tightly controlled. California generally prohibits re-soliciting consent for the same high-risk AI use for 12 months, while Colorado permits re-consent only through a neutral and accessible interface accompanied by detailed disclosures explaining how the AI system works, how it affects decisions, and the potential consequences of renewed use.

Procedural Rules Add Operational Complexity

States also regulate *how* consumers may exercise access and opt out rights.

Most laws require multiple, accessible submission methods aligned with the business's primary modes of interaction with the consumer. They also prescribe how businesses must respond, including confirmation that a request was received and processed, standardized timelines for compliance, and clear response formats designed to be understandable to consumers.

Verification rules apply to most rights, but some laws limit identity verification for opt out requests, reflecting the view that opting out should be friction-free.

What Consumer Control Means for Businesses

Together, these obligations shift AI governance toward outcome-based accountability, where consumer rights can directly determine whether and how AI systems may be used to make significant decisions.

Businesses deploying high-risk AI systems must understand, document, and clearly explain how these AI systems function in practice, including how personal information is processed, how outputs influence decisions, and when human judgment meaningfully intervenes. This places operational pressure on businesses to align technical system design, internal governance, and consumer facing explanations, and to maintain processes that allow decisions to be reevaluated when data is corrected or challenged.

Opt out rights and related procedural requirements similarly require businesses to build enforceable controls over AI use, not just theoretical consumer choices. Companies must be able to stop or adjust high-risk AI processing within prescribed timelines, honor opt out requests across systems and downstream service providers, and carefully manage how and when consent may be reintroduced.

PART 5

Assessing Risk Before AI Decides

In the last two installments of **Legal Lines Around AI**, we examined how transparency and consumer rights work together to give individuals greater control over how businesses use high-risk AI to make consequential decisions about them.

Those obligations are reinforced and operationalized through risk assessments, which have quickly become a centerpiece of AI governance. Because high-risk AI systems can amplify risks such as bias, unfair treatment, and invasive profiling, state laws increasingly require businesses to assess those risks before deployment, document safeguards, and revisit their analysis as systems and use cases evolve.

State Laws Converge on Pre-Deployment Risk Assessments

States are converging on a core rule: deployers of high-risk AI systems must complete a documented risk assessment before deployment. This obligation applies regardless of whether the system is developed internally or obtained from a vendor or third-party provider.

California requires risk assessments before processing personal information that presents significant risk to privacy, including inputting personal information in high-risk AI systems, automated employment-related inferences, inferences based on sensitive locations, and the training of high-risk or biometric technologies.

Colorado requires deployers to complete a risk assessment before deploying *any* high-risk AI system. Other state comprehensive privacy laws similarly mandate documented assessments when automated processing creates reasonably foreseeable risks, such as unfair or deceptive treatment, unlawful disparate impact, financial or physical injury, offensive intrusions into private affairs, or other substantial consumer harm.

Across these laws, the unifying principle is risk-based. Regulators are less concerned with labels and more focused on whether and how high-risk AI use meaningfully affects individuals' rights or wellbeing.

What Regulators Expect in a High-Risk AI Risk Assessment

A compliant high-risk AI risk assessment should tell a complete and defensible story: how the high-risk AI system works, the risks it creates, and why deploying it is justified given those risks.

At a minimum, assessments must:

- Clearly define the specific purpose of the high-risk AI processing
- Describe the processing activity itself and identify the categories of personal information involved
- Explain the context of the processing, the organization's relationship with affected consumers, and consumers' reasonable expectations

Intended use cases must closely align with the stated purpose, with particular focus on whether the AI system is used to make or materially influence significant decisions.

Beyond these fundamentals, assessments are expected to address operational detail, risk analysis, and mitigation. This includes:

- How data is collected, used, retained, and disclosed
- What transparency measures and disclosures are provided
- How outputs are used in decision-making
- The role of human oversight
- Whether third-party tools or vendors are involved

Regulators also expect a thorough evaluation of foreseeable harms, including privacy, discrimination, financial, psychological, and constitutional risks, along with a documented explanation of how safeguards reduce those risks and why the benefits outweigh remaining concerns.

In practical terms, regulators expect risk assessments to address:

- Purpose and scope of AI-driven processing, with specificity
- Data inputs and outputs, including sensitive or children's data
- How the system works, including logic, assumptions, limitations, and training data
- Risks, such as discrimination, unfair treatment, privacy intrusion, or economic harm
- Safeguards and mitigation measures, including security, governance and bias controls
- Transparency and oversight mechanisms, including consumer notices, monitoring, audits, and metrics
- Decision-making accountability, identifying who approved deployment and why

Stakeholder involvement is mandatory: employees involved in the processing activity must also participate in AI risk assessments by providing operational details such as how data is collected, used, and managed to ensure the assessment reflects real world system use.

Risk Assessments as Living Documents

Finally, assessments should reflect an ongoing process. Regulators increasingly expect evidence of post-deployment monitoring, regular review, internal or external audits, and clear ownership within the organization.

High-risk AI risk assessments are not one-time paperwork; they must be reviewed and updated throughout the system's lifecycle. At a minimum, assessments should be revisited annually. More importantly, deployers must update assessments whenever there is a material change to the processing activity that could create new risks, increase existing risks, or weaken existing safeguards (i.e., changes to data types or sources, processing purpose, algorithms, vendors, software, or system outputs). These updates must be completed as soon as feasible and no later than 45 days after the material change.

Retention and regulatory access are equally critical. Deployers must retain all versions of risk assessments (original and updated) for as long as the processing continues and for five years after completion, whichever is later. Organizations must also be prepared to produce these assessments to regulators on demand, typically within 30 days.

In California, deployers face an added layer of oversight: certain risk-assessment details must be formally submitted to the California Privacy Protection Agency on a defined schedule, accompanied by an executive attestation signed under penalty of perjury.

Practical Considerations for Deployers of High-Risk AI

As AI risk assessment requirements shift from theory to enforcement, businesses should focus on building processes that are durable, defensible, and scalable.

Deployers should consider the following steps:

- **Inventory AI systems early and often.** Maintain an up-to-date inventory of AI tools used for decision-making, profiling, or inference, including vendor-provided systems and internally developed tools.
- **Align AI risk assessments with existing privacy governance.** Integrate AI risk assessments into existing DPA or privacy risk assessment workflows rather than creating siloed processes.
- **Engage operational stakeholders at the outset.** Involve teams responsible for data collection, model operation, and business use cases to ensure assessments reflect actual system behavior, not theoretical design.
- **Plan for change management.** Establish triggers and internal processes to identify material changes to AI systems or data use and ensure assessments are updated within required timelines.
- **Document human oversight and escalation paths.** Be explicit about when and how humans review AI outputs, override decisions, or handle consumer challenges.
- **Prepare for regulator access.** Retain assessments in a centralized, review-ready format and assign ownership so the organization can respond quickly to regulatory requests.
- **Do not rely solely on vendor assurances.** Even when AI tools are third-party, deployers remain accountable for understanding risks, evaluating safeguards, and documenting compliance.

In practice, regulators are less focused on whether a document exists and more so on whether the assessment reflects real operational decision-making and ongoing oversight. Businesses that invest now in practical, repeatable assessment processes will be far better positioned to adapt to evolving legal expectations.

PART 6

Building an AI Governance Program That Scales

Across the United States, AI regulation is evolving quickly but not randomly. While state laws vary in scope, applicability, and mechanics, they are converging on a shared principle: AI systems that meaningfully affect people must be governed through risk-based oversight, transparency, and accountability.

For businesses operating nationally, the challenge is not mastering a single statute. It is building an AI governance program that can absorb regulatory change while remaining practical to implement. In this final installment of **Legal Lines Around AI**, we bring the series together by focusing on how organizations can design an AI compliance posture that meets the toughest requirements without becoming unworkable.

Design for the Most Demanding Laws—Once

The most resilient AI governance programs are built around the highest regulatory expectations, not the lowest common denominator.

Several states already require pre-deployment risk assessments, plain-language disclosures, consumer control mechanisms, documented safeguards, and ongoing oversight for high-risk or consequential AI systems. California and Colorado currently impose the most comprehensive and demanding requirements. A governance program designed to satisfy those standards will generally exceed obligations elsewhere and remain durable as new laws are enacted.

This approach allows organizations to implement one core governance architecture, layered with limited state-specific adjustments where necessary, rather than managing fragmented compliance workflows.

Classification Is the Real Trigger Point

Nearly every obligation discussed throughout this series flows from a single determination: what kind of AI system is being used, and how is it used?

Strong governance requires early, repeatable classification, especially whether a system qualifies as "high risk" or is used to make or substantially influence significant decisions about individuals. That determination should not live solely with legal teams. It needs to be embedded in product intake, procurement, and deployment processes, and revisited whenever a system's purpose, data inputs, or outputs materially change.

Without disciplined classification, even well-intentioned governance programs will struggle to apply the right controls at the right time.

Make Risk Assessments the Center of Gravity

Risk assessments are the most effective way to operationalize AI governance.

Scalable programs treat AI risk assessments as living documents that follow a system throughout its lifecycle. Rather than creating one-off assessments for every use case, mature programs reuse and consolidate assessments where risks and systems are comparable, supplementing them only when necessary. This approach is explicitly permitted under several state laws and is essential for scale.

Well-designed assessments do more than satisfy regulators. They create internal clarity around system purpose, data use, decision logic, human involvement, foreseeable harms, and mitigation measures. When assessments are tied directly to deployment approval and ongoing monitoring, governance becomes part of normal business operations—not an after-the-fact compliance exercise.

Transparency Should Match the Consumer Experience

AI transparency obligations increasingly apply at multiple points, including:

- **Public disclosures** (e.g., websites or AI use-case inventories)
- **Pre-use notices** before consequential decisions
- **Post-decision explanations** and appeal rights

Effective governance aligns legal disclosures with product design and customer communications. Plain-language explanations, consistent terminology, and centralized disclosure templates reduce risk while building trust. Accessibility is no longer optional and should be baked into disclosure workflows from the outset.

When transparency is treated as a user experience issue rather than a legal afterthought, compliance is easier to maintain and explain at scale.

Human Oversight Must Be Real, Not Symbolic

AI laws provide varying flexibility where meaningful human involvement exists, particularly with respect to opt-out rights and adverse decisions. But regulators are increasingly skeptical of vague or symbolic claims of human review.

Strong governance programs clearly document:

- When humans review AI outputs
- Who has authority to override decisions
- How appeals and challenges are handled
- Training provided to reviewers
- Escalation paths for bias, error, or system failure

These processes must be consistently applied and auditable. Simply stating that "a human is involved" without operational detail may cause regulators to treat a system as fully automated anyway.

Integrate AI Governance into Existing Compliance Structures

The strongest AI governance programs do not stand alone. Instead, they integrate AI oversight into familiar compliance frameworks, including:

- Privacy impact and data protection assessments
- Vendor risk management
- Product intake and launch reviews
- Security and data governance programs
- Incident response and escalation processes

AI governance works best when it feels familiar, even if the technology is new. Organizations that leverage existing controls, committees, and workflows can scale oversight without reinventing their compliance infrastructure.

Accountability Must Be Clear and Executive-Level

AI laws increasingly expect clear ownership and executive accountability. Risk assessments, public disclosures, and regulatory submissions often require sign-off from individuals with authority.

Effective governance programs define:

- AI governance owners
- Cross-functional participants
- Executive approvers
- Documentation and retention responsibilities

This clarity enables faster, better decisions, especially when tradeoffs between innovation and risk must be resolved.

Closing Thought: Build for Change, Not Certainty

U.S. AI regulation will continue to evolve, but its direction is already clear. The goal is not to predict every future law, it is to build a governance posture that can adapt without constant reinvention.

Organizations that anchor their programs in risk-based assessments, meaningful transparency, real human oversight, and integrated compliance structures will be best positioned to scale AI responsibly and remain compliant as legal expectations continue to rise.

If you need assistance with or have questions about the topics covered in this series, please reach out to:

Aaron Parry

aparry@mslawgroup.com

(614) 939-9955