# Countdown to Compliance:

## The Comprehensive CEO & CIO Guide to the Colorado AI Act


September 2025

GRANITE FORT
ADVISORY

Granite Fort Advisory
**AI Transformation, Governance, Risk & Compliance**
Clarity. Compliance. Confidence.

www.granitefort.com

GRANITE FORT ADVISORY

# Executive eBook

## CONTENTS

# EXECUTIVE SUMMARY

This eBook provides an in-depth, under-the-hood analysis of the Colorado AI Act, offering a granular look at its implications for organizations. For a quick review, please see this Executive Summary section or the companion **PowerPoint slide deck**.

## Why the Colorado AI Act Matters

In May 2024, Colorado became the first U.S. state to enact a comprehensive AI accountability law. It represents a paradigm shift from voluntary guidelines to enforceable governance.

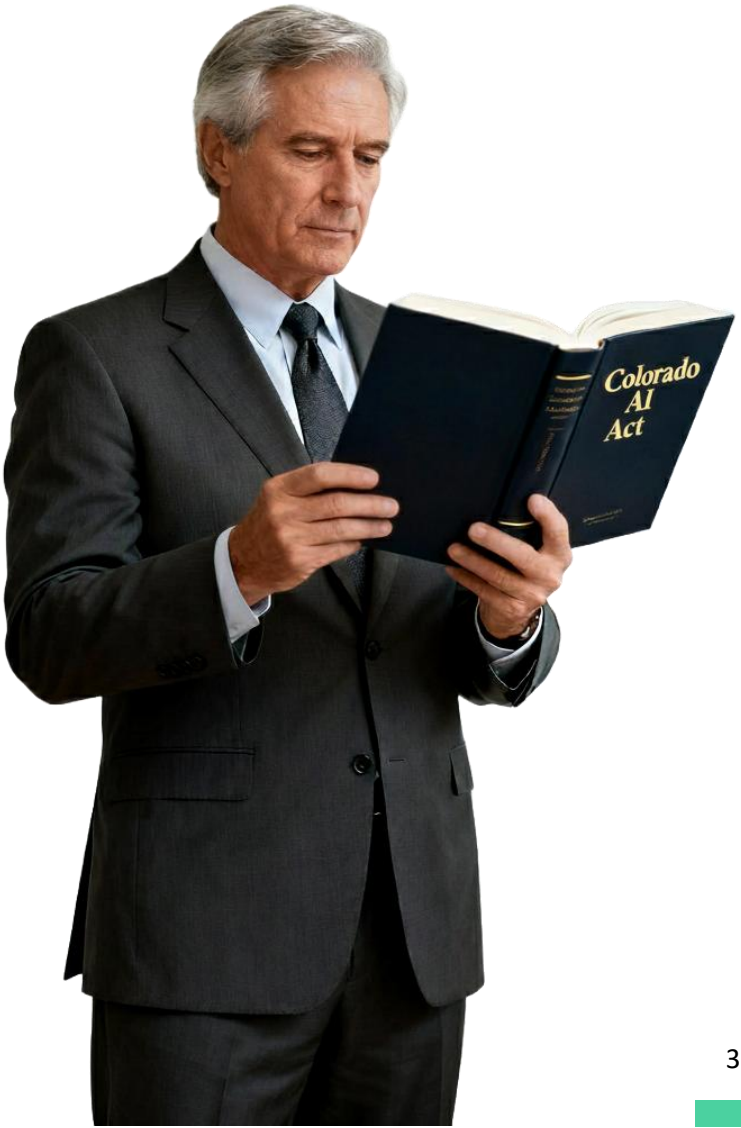


Formally known as Senate Bill 24-205, the Colorado Artificial Intelligence Act (CO AI Act or the Act) creates legal duties for businesses that develop or deploy "high-risk AI systems." Enforcement begins in June 2026, giving companies limited time to prepare.

For CEOs and CIOs, this is not a niche regulatory issue. It is a strategic inflection point:

- AI is now regulated like financial services, data privacy, and consumer protection.
- Compliance failures could cost millions in fines and reputational damage.
- Proactive governance will differentiate trusted enterprises from laggards.

## Why This Law Is Different

For any organization with business operations in Colorado, this regulation demands immediate attention from the highest levels of leadership. This landmark piece of legislation is not merely a local compliance concern; it is a bellwether for what is to come across the nation, establishing a risk-based framework that mirrors the key provisions of the European Union's AI Act.


Unlike prior data privacy or consumer protection rules, the Act targets the black box of AI decision-making. The law specifically addresses algorithmic discrimination - a risk that has drawn intense regulatory, political, and media scrutiny.

## Executives must recognize three shifts:

1. **AI as Regulated Infrastructure**: High-risk AI is no longer a purely technical tool. It is regulated infrastructure, like banking systems or medical devices.
2. **Board Accountability**: CEOs, CIOs, and boards cannot delegate AI compliance to the back-office "data science team." The Act requires enterprise-wide governance with monitoring at the highest levels.
3. **Trust as a Strategic Asset**: Transparent, fair AI will become a market differentiator. Enterprises that can prove fairness and accountability will build stronger consumer trust.

## Risks of Non-Compliance

- **Financial**: Each discriminatory outcome could be a separate $20,000 violation. A single flawed algorithm affecting thousands of consumers could result in $ Millions in exposure.
- **Operational**: The Attorney General can demand audits, impose injunctions halting AI use, or require system redesigns - disrupting critical operations.
- **Reputational**: Public perception is paramount. Headlines about "biased AI" damaging vulnerable populations can erode consumer and investor trust overnight.
- **Strategic**: Competitors that align early with the Act's standards may gain procurement advantages, attract investors and reduce compliance costs.

## Key Takeaways for Senior Leadership

The most important takeaway for senior leadership is that the recent five-month delay in the Act's effective date - from February 1, 2026, to June 30, 2026 - is not a license to pause preparations. Instead, this reprieve presents a crucial opportunity to accelerate. **The most effective strategy is a proactive one**: leverage this time to adopt a globally AI recognized governance framework such as ISO/IEC 42001, and begin building the required documentation and internal processes. This approach not only prepares the organization for the current law but also builds a "rebuttable presumption" of compliance, offering a powerful legal defense and positioning the company to navigate a dynamic regulatory environment that will likely see future state-level and federal regulations.

| Issue | Executive Impact |
|---|---|
| **Scope** | Applies to most medium-to-large organizations doing business in Colorado, including national firms with Colorado customers. |
| **Obligations** | Both **developers** (creators) and **deployers** (users) of high-risk AI must document risks, mitigate discrimination, and provide transparency. |
| **Enforcement and Penalties** | Violations = unfair or deceptive trade practices, enforceable by Attorney Fines up to **$20,000 per violation**. Violations may also expose deployers to private lawsuits under the Colorado Consumer Protection Act (including treble damages and attorney fees) and federal civil rights claims, potentially resulting in **multimillion-dollar exposure** from class actions. |
| **Timeline** | Compliance deadline: **June 30, 2026** (delayed from original Feb 2026). |
| **Safe Harbors** | Adoption of standards like **ISO/IEC 42001** offers rebuttable presumption of compliance. |
| **Strategic Importance** | AI governance is now a **board-level responsibility**. Early movers will win consumer trust and investor confidence. |

## Bottom Line for Senior Leaders

The Colorado AI Act is a boardroom issue, not just a compliance checklist. CEOs and CIOs must:

1. Treat AI like a regulated asset.
2. Establish AI governance programs aligned with global frameworks.
3. Prepare disclosures, documentation, and impact assessments now.
4. Monitor legislative updates - but assume compliance will be required by mid-2026.
5. Expect this law, or portions thereof, to be replicated by other states nationwide

Failure to act risks financial penalties and reputational harm. Success positions your enterprise as a leader in trustworthy AI.

The Colorado AI Act isn't just another regulatory burden - it's a catalyst for establishing world-class AI governance that will serve your organization well beyond Colorado's borders. Organizations that view this as an opportunity rather than an obligation will emerge as leaders in the responsible AI economy.

# KEY DEFINITIONS UNDER THE ACT

The Colorado AI Act introduces several key definitions that are critical for understanding scope and obligations.

## What is an "Artificial Intelligence System"?

The Colorado AI Act begins with a broad and inclusive definition of an artificial intelligence system. The law defines an artificial intelligence system as "*any machine-based system that, for any explicit or implicit objective, infers from the inputs the system receives how to generate outputs, including content, decisions, predictions, or recommendations, that can influence physical or virtual environments*".

This definition is intended to be comprehensive, encompassing everything from simple decision-making algorithms to complex generative models. It is the foundation upon which all other key definitions and obligations are built.

## What are "High-Risk AI Systems"?

The primary focus of the Act is on "high-risk artificial intelligence systems." An AI system is classified as high-risk if it "*makes, or is a substantial factor in making, a consequential decision*". Consequential decisions are those with material legal or similarly significant effects on consumers, including access to or denial of services in 8 enumerated areas: education, employment, financial or lending services, essential government services, healthcare services, housing, insurance and legal services. Examples include AI-driven resume screening tools in hiring, credit scoring algorithms in lending, or diagnostic aids in healthcare.

Not all AI falls under this umbrella - for example, low-risk chatbots for customer service, for instance, are exempt unless they influence consequential outcomes. This risk-based approach is central to the law's design, concentrating regulatory efforts on the most impactful uses of AI. The law provides specific carve-outs for certain technologies, excluding a wide range of systems that perform narrow procedural tasks or do not directly influence consequential decisions.


These include, but are not limited to, anti-malware, anti-virus software, AI-enabled video games, calculators, and spam filters. However, a critical caveat exists: these technologies are only exempt "unless the technologies, when deployed, make, or are a substantial factor in making, a consequential decision". This condition means that a seemingly benign system could, depending on its specific application, fall under the high-risk category and trigger all of the Act's obligations.

## The Nuances of "Algorithmic Discrimination" and "Disparate Impact"

The core harm the Act seeks to prevent is algorithmic discrimination. The law defines this as "*any condition in which the use of an artificial intelligence system results in an unlawful differential treatment or impact that disfavors an individual or group of individuals on the basis of their actual or perceived protected class*". This mirrors anti-discrimination principles in civil rights law but applies specifically to AI outputs.

The phrase "*differential... impact*" is of paramount importance because it explicitly includes the concept of "disparate impact." Unlike some discrimination laws that require proof of intent to discriminate, this Act holds a company accountable for the unintended, discriminatory outcomes of its AI systems.

In other words, algorithmic discrimination is not limited to intentional bias; unintentional disparate impacts from biased training data or flawed algorithms also qualify. A system can be designed with no malicious intent, yet if its use results in a disproportionately negative outcome for a protected group, it could be considered a form of algorithmic discrimination. This shift in legal exposure from intent to outcome is the most significant strategic point for executives, as it requires a fundamental change in how AI systems are designed, tested, and deployed. Executives must recognize that even subtle biases in AI can trigger liability, emphasizing the need for rigorous testing and mitigation.

## Defining the Roles: "Developer" vs. "Deployer"

The Act creates a bipartite liability framework, assigning distinct responsibilities to two key roles within the AI ecosystem.

A "**developer**" is defined as a person doing business in Colorado that "develops or intentionally and substantially modifies" an AI system. This includes any deliberate change that results in a new, reasonably foreseeable risk of algorithmic discrimination. The definition explicitly excludes changes that result from a system's continuous learning after deployment.

Thus, developers are persons or entities that develop or substantially modify an artificial intelligence system including:

- Original creators of AI systems
- Organizations that significantly customize pre-existing systems
- Vendors providing AI systems to other organizations
- Internal teams building proprietary AI solutions

A "**deployer**" is a person doing business in Colorado that deploys a high-risk AI system. Deployers are persons or entities that use high-risk AI systems. This includes:

- Organizations using third-party AI tools in their operations
- Companies integrating AI into customer-facing services
- Entities using AI for internal decision-making processes
- Any organization whose AI use affects Colorado residents.

The law is structured to allocate responsibility and compliance burdens based on these roles, recognizing the different points of control and influence in the AI supply chain.

Many organizations will be both developers and deployers, triggering dual compliance obligations. For example, if your company builds a proprietary hiring algorithm (Developer) and uses it to screen candidates (Deployer), you must satisfy requirements for both roles.

## "Consequential Decisions" and "Substantial Factor"

The definition of a "**consequential decision**" is the linchpin that determines if an AI system is "high-risk." The law defines a consequential decision as one that has a "material legal or similarly significant effect" on a person's life in one of eight key areas: education, employment, financial or lending services, essential government services, healthcare services, housing, insurance, or legal services.

A "**substantial factor**" is any use of an AI system to generate content, a decision, a prediction, or a recommendation concerning a consumer that is used "as a basis to make" a consequential decision. This definition means that even an AI system that doesn't make a final decision but provides a critical recommendation or prediction can trigger the law's requirements. For example, a résumé-screening tool that ranks candidates could be a "substantial factor" in an employment decision, even if a human makes the final hiring choice. This pushes the boundaries of accountability to include systems that influence, rather than just make, a consequential decision.

To provide a clear, at-a-glance reference for senior leaders, the following table summarizes key definitions and their applicability criteria.

| Key Definition | Explanation | Applicability Criteria |
|---|---|---|
| **Artificial Intelligence System** | A machine-based system that infers from inputs to generate outputs that can influence physical or virtual environments. | This is the base definition; the Act's obligations apply to specific types of AI systems. |
| **High-Risk AI System** | An AI system that makes, or is a substantial factor in making, a consequential decision. | The law's core obligations apply only to these systems. Excludes procedural tools unless they become a substantial factor in a consequential decision. |

| | | |
|---|---|---|
| **Algorithmic Discrimination** | Any condition where an AI system's use results in unlawful differential treatment or impact that disfavors a person based on a protected class. | The core harm the law aims to prevent. Explicitly includes unintentional "disparate impact." |
| **Developer** | A person doing business in Colorado that develops or intentionally and substantially modifies an AI system. | Developers must fulfill documentation, public disclosure, and risk reporting obligations. |
| **Deployer** | A person doing business in Colorado that uses a high-risk AI system. | Deployers must fulfill risk management, impact assessment, and consumer notification obligations. |
| **Consequential Decision** | A decision with a material legal or similarly significant effect on a consumer's access to or cost of a service in one of the 8 key areas like education, employment, financial or lending services, essential government services, healthcare services, housing, insurance, or legal services. | If an AI system influences a decision in one of these areas, it is likely a high-risk system. |

## Impact Assessment: Colorado AI Act's Pillar for Safe AI

Impact Assessment has been given a lot of importance and serves as a cornerstone of the Act, ensuring that every high-risk AI deployment undergoes rigorous review, risk mitigation, transparency, and ongoing oversight to protect consumers and uphold fairness.

Impact Assessment, as defined in Section 6-1-1703(3)(b) of the Act, is a documented review that a deployer of a high-risk artificial intelligence system must complete and that, to the extent reasonably known by or available to the deployer, must include at a minimum:


- a statement disclosing the purpose, intended use cases, deployment context of, and benefits afforded by the high-risk artificial intelligence system;
- an analysis of whether deployment poses any known or reasonably foreseeable risks of algorithmic discrimination and, if so, the nature of that discrimination and the steps taken to mitigate the risks;
- a description of the categories of data the system processes as inputs and the outputs it produces;
- if the deployer used data to customize the system, an overview of the categories of data used for that customization;
- any metrics used to evaluate performance and known limitations of the system;
- a description of any transparency measures taken, including consumer disclosures when the system is in use; and
- a description of post-deployment monitoring and user safeguards, including the oversight, use, and learning processes established to address issues arising from deployment.

## Risk Management under the Act

At the core of the Colorado AI Act, Risk Management Policy and Program is the mandatory, documented framework that deployers of high-risk AI systems must implement to systematically identify, document, and mitigate algorithmic discrimination risks. It must specify the principles, processes, and personnel used to identify, document, and mitigate known or reasonably foreseeable risks of algorithmic discrimination.


It must be planned, implemented, and regularly reviewed and updated over the system's life cycle and align with ISO/IEC 42001, the NIST AI RMF or another nationally or internationally recognized risk management standard deemed substantially equivalent or more stringent; the deployer's size and complexity; the nature, scope, and intended uses of its high-risk AI systems; and the sensitivity and volume of data processed.

A single risk management policy and program may cover multiple high-risk AI systems.

# SCOPE AND APPLICABILITY

## Who is Covered? The "Doing Business in Colorado" Criterion

The Colorado AI Act's jurisdiction is broad, extending beyond state borders. The law applies to both developers and deployers of high-risk AI systems that are "doing business in Colorado". This means that companies headquartered outside of Colorado that sell, license or deploy high-risk AI systems to entities operating within the state are also subject to the law's requirements. This **extraterritorial reach is critical for national and international businesses to understand**, as it necessitates a comprehensive review of all their AI systems that could be used by Colorado-based customers or employees.

## The Small Deployer Exemption and Other Carve-Outs

The Act provides a limited exemption for certain small deployers. A deployer is exempt from the most burdensome requirements if they satisfy three criteria: they have fewer than 50 employees, they do not use their own data to train the AI system, and they make any impact assessment completed by the developer available to consumers. This exemption is not an automatic shield; it is conditional.

A small business relying on a third-party AI system still has a legal obligation to ensure their vendor provides the necessary documentation to satisfy the third criterion.

Beyond the small deployer exemption, the Act also carves out a long list of specific technologies from the "high-risk" definition, including anti-virus software, calculators, video games, and spam filters. However, as noted previously, this exclusion is conditional.

A seemingly benign system is not immune if it is used to make or is a substantial factor in making a consequential decision. For instance, a chatbot designed to answer simple questions could unexpectedly become a substantial factor in a consequential decision if it provides a financial recommendation that is used as the basis for a lending decision. This creates a regulatory trap for the unwary and underscores the necessity of a thorough internal audit of all AI use cases.

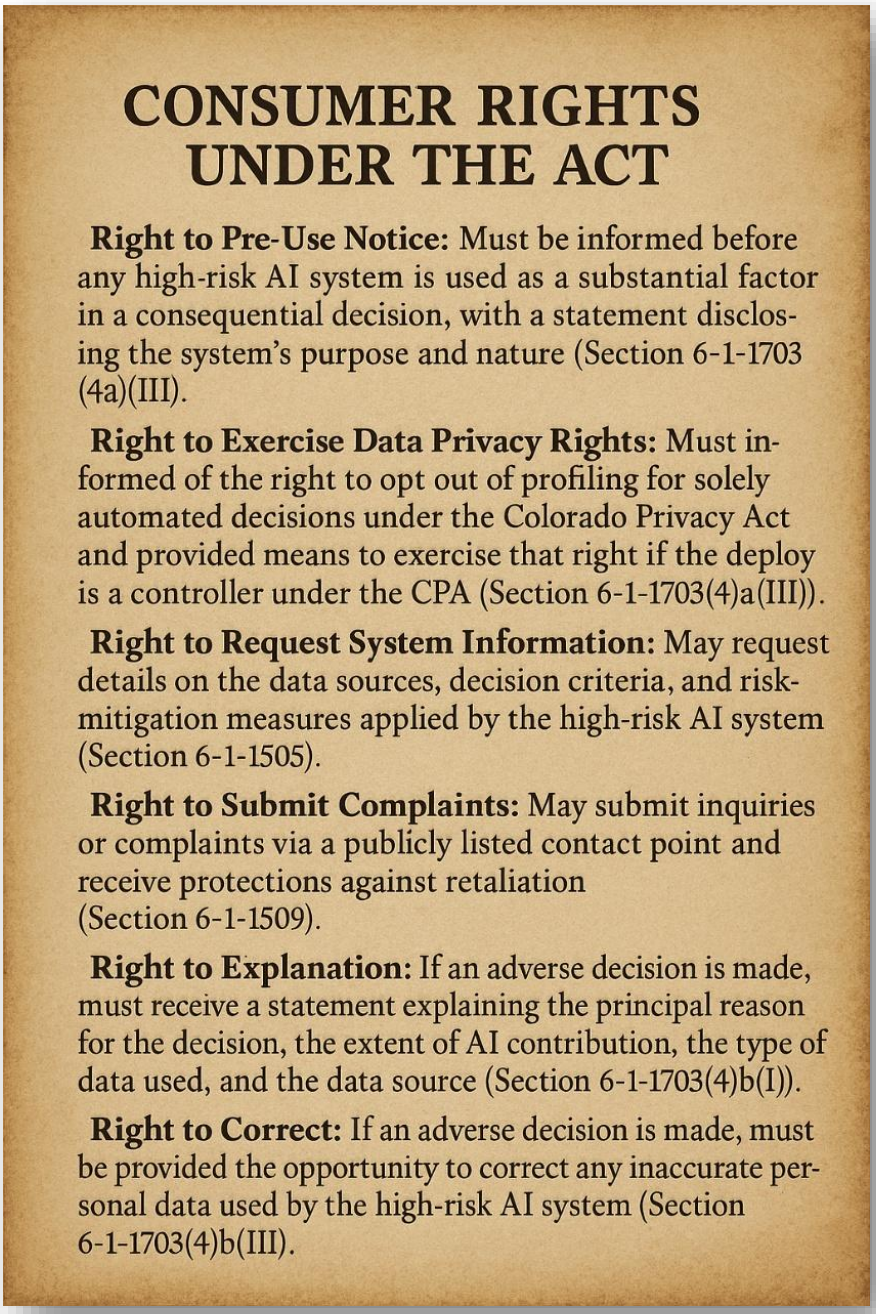


## The "High-Risk" Filter: Systems to Inventory and Assess

The first and most crucial step for any organization is to conduct a comprehensive inventory of all AI systems currently in use or planned for future deployment. This requires mapping each system's function against the Act's definition of "consequential decisions."

Businesses must ask:

- Does this AI system play a role in hiring, firing, or promotions?
- Is it used to evaluate loan applications, insurance policies or healthcare decisions, etc.?
- Does it affect access to housing or essential government services?

A failure to identify and assess every system that falls under this definition exposes the organization to significant legal and financial risk. The law's design pushes companies to identify their exposure proactively, rather than waiting for an enforcement action.

## Navigating Cross-Jurisdictional Challenges

The Colorado AI Act's extraterritorial reach, coupled with the absence of a federal standard, creates a fragmented and complex regulatory landscape for companies operating nationally. The law sets a unique and demanding benchmark for AI governance, particularly with its focus on disparate impact, which currently contrasts with the federal government's enforcement posture.

For national companies, this "**patchwork**" of state-by-state regulations can be a logistical and compliance nightmare. A reactive, state-specific compliance strategy is inefficient and can lead to a fragmented governance model.

The most strategic approach is to design and implement a harmonized, enterprise-wide AI governance program that can meet the most stringent requirements with Colorado's law serving as the de facto benchmark. **By meeting Colorado's high bar, a company will be well-positioned to comply with emerging regulations in other states**.

# EVOLVING TIMELINE OF THE **ACT**

## The Initial Landmark: Passage and First Effective Date

The Colorado AI Act, Senate Bill 24-205, was signed into law by Governor Jared Polis on May 17, 2024. This was a landmark moment, making Colorado the first state in the U.S. to enact a comprehensive AI regulatory framework. The law was originally slated to become effective on February 1, 2026, which provided organizations with a specific, albeit tight, timeline to prepare for compliance.

## A Political Reprieve: The Five-Month Delay to June 30, 2026

The initial effective date was met with significant industry pushback and concerns from the Governor himself regarding the law's complexity and its potential to stifle innovation. Following a special legislative session held in late August 2025, a new bill, SB 25B-004, was passed and signed into law. This bill delayed the effective date of the Act by five months, pushing the principal operative dates back to June 30, 2026.

## Legislative Intent and Unforeseen Consequences

The legislative history behind this delay reveals deep divisions and failed compromises. Lawmakers had attempted to pass substantive amendments to the Act to address issues like the broad definition of "algorithmic discrimination," the scope of exemptions, and the protection of trade secrets. However, these efforts collapsed due to intense industry lobbying and a lack of consensus. The five-month delay was ultimately a procedural solution to buy more time for a potential overhaul in the next regular legislative session, not a sign of the law being abandoned. It was a political compromise to address the complexity of implementation without fully retreating from the law’s core principles.

## Strategic Implications of the Delay: A Time to Act, Not Wait

The five-month delay offers a critical window for businesses to prepare. However, taking a "wait-and-see" approach, as some have suggested, is a high-risk gamble. The failed attempts to "pare

down" the law indicate that a political consensus on a less stringent version has not yet been reached. This suggests that the current, demanding framework is the most likely starting point for compliance. A forward-looking strategy dictates that organizations use this time to accelerate their preparations: conduct comprehensive AI inventories, draft policies, and pilot compliance frameworks. This proactive approach reduces risk, builds a strong foundation for a future-proof governance program, and positions the company to act quickly if the law is not significantly altered.

The most nuanced understanding of this situation is that the delay is not a sign of the law's demise but a strong indicator that its core principles - the duty of care, impact assessments, and transparency - are likely to be refined rather than abandoned. Waiting for a more business-friendly version to pass is a risky strategy. The prudent course of action is to get a head start on the current requirements, knowing that any future changes will likely make compliance easier, not harder.

The following table provides a clear, at-a-glance representation of the key dates in the Colorado AI Act's legislative timeline:

| Event | Date | Significance |
|---|---|---|
| **Passage of SB 24-205** | May 17, 2024 | Colorado becomes the first U.S. state to pass comprehensive AI regulation. |
| **Original Effective Date** | February 1, 2026 | Initial deadline for compliance with all requirements. |
| **Passage of SB 25B-004** | August 28, 2025 | Bill delaying the effective date is signed into law. |
| **New Effective Date** | June 30, 2026 | New deadline for compliance, providing a five-month reprieve for businesses. |

# OBLIGATIONS FOR DEVELOPERS (AI VENDORS)

In simple language, the obligations could be summarized as Duty of Care and Transparency.

## The "Reasonable Care" Standard

On and after June 30, 2026, a developer of a high-risk AI system must use "reasonable care" to protect consumers from any "known or reasonably foreseeable risks of algorithmic discrimination" arising from the system's intended and contracted uses. This is a critical distinction, as the law does not impose strict liability. Compliance will be assessed in light of recognized frameworks (such as ISO/IEC 42001) and contextual factors for deployers (i.e. intended use cases, complexity, sensitivity of data, scope of the high-risk AI, etc. for their corporate Customers). This reflects a context-sensitive evaluation to determine if developers exercised due care.

## The Documentation Imperative: What to Provide to Deployers

A central obligation for developers is the provision of a "packet" of documentation to deployers. This information is crucial for enabling deployers to meet their own compliance requirements. This packet must include:

- A general statement describing the system's reasonably foreseeable uses and any known harmful or inappropriate uses.
- Detailed documentation on the type of data used to train the system, its purpose, intended benefits, and known limitations, including risks of algorithmic discrimination.
- Documentation describing how the system was evaluated for performance and bias mitigation, the data governance measures used, and how it should be used, not be used, and monitored by an individual when making a consequential decision.

The Act suggests using industry standard artifacts such as "model cards," "dataset cards," or other impact assessments to provide this information. The law's design makes developers the front-line data providers for the AI ecosystem in Colorado, placing a clear legal and contractual imperative on them to support their customers' compliance efforts. A deployer cannot be compliant without a compliant developer, effectively extending the law's influence far beyond Colorado's borders to any developer who wishes to sell to businesses operating there.

## Public-Facing Disclosures and Website Statements

Developers must maintain a clear, regularly updated public statement on their website or in a public use case inventory. This public-facing disclosure must summarize the types of high-risk AI systems they make available and how they manage the known or reasonably foreseeable risks of algorithmic discrimination. This requirement is intended to foster public transparency and accountability.

## Reporting to the Attorney General: The 90-Day Rule

Within 90 days of discovering, or receiving a credible report, that their system has caused or is likely to cause algorithmic discrimination, developers must disclose this information to the Colorado Attorney General and all known deployers of the system. This is a mandatory and proactive reporting obligation that places a clear duty on developers to monitor their systems for discriminatory outcomes and to inform their customers and the regulator if such issues arise.

## Managing Continuous Learning and Substantial Modifications

The duty of care is not a one-time event; it is an ongoing obligation. An "intentional and substantial modification" is defined as a deliberate change to an AI system that results in any "new reasonably foreseeable risk of algorithmic discrimination". This distinction is important for continuously learning systems. The Act clarifies that a change resulting from a system's continuous learning after deployment is not considered an intentional and substantial modification, thereby avoiding a constant cycle of re-evaluation for models that are designed to evolve post-deployment.

# OBLIGATIONS FOR DEPLOYERS (COMPANIES USING AI)

In simple language, the obligations could be summarized as Governance, Assessments and Consumer Rights.

## Establishing an AI Risk Management Policy and Program

Just as with developers, deployers are subject to a "reasonable care" standard. To meet this standard, a deployer must establish and maintain a “risk management policy and program”. This program must describe the principles, processes, and personnel used to identify, document, and mitigate known or reasonably foreseeable risks of algorithmic discrimination. The law provides a significant incentive for compliance by granting a rebuttable presumption of reasonable care if a deployer can demonstrate that they have implemented such a program.

## The Mandate for Impact Assessments: Annual and Triggered Reviews

The law requires deployers to conduct an impact assessment annually and within 90 days after making an intentional and substantial modification to a high-risk AI system. This assessment must be comprehensive, including a risk analysis of potential algorithmic discrimination, an analysis of the system's data and outputs, and a description of the metrics used for performance evaluation. This recurring obligation means that a deployer's job is never truly finished; they must maintain a continuous monitoring and assessment program to ensure ongoing compliance.

## Consumer Notification Requirements: Pre-Decision and Post-Adverse-Decision

Deployers must provide clear and direct notifications to consumers. When a high-risk AI system makes or is a substantial factor in a consequential decision, the deployer must notify the consumer at or before the decision is made. If the decision is adverse to the consumer, a more detailed notification is required. This notification must disclose the principal reasons for the decision, the degree to which the AI system contributed to it, and the types and sources of data that were processed. Additionally, for any AI a consumer interacts with (not just high-risk ones), if it's not obvious that they are interacting with an AI, the deployer must disclose this fact.

## Fulfilling Consumer Rights: Data Correction and Human Review Appeal

The Act grants consumers two key rights related to adverse consequential decisions. **First**, consumers have the right to correct any incorrect personal data that was processed by the AI system to make the decision. **Second**, they have the right to appeal an adverse decision, with the opportunity for human review if technically feasible. The right to appeal with human review creates a direct, real-time channel for consumers to flag potential issues. If a business receives multiple appeals citing the same issue, it is a strong indicator of a systemic problem, which would then trigger the need to re-evaluate the system and potentially update the impact assessment. Operationalizing these consumer rights is a significant and ongoing business expense.

## Deployer Website Transparency: Managing Public Perception

Deployers must also maintain a publicly available statement on their website. This statement should summarize the types of high-risk AI systems they currently deploy, how they manage known or foreseeable risks of algorithmic discrimination, and the nature, source, and extent of the information collected and used by the system. This public-facing transparency is a central theme of the law and is intended to build consumer trust and accountability.

## The Vendor Management Imperative

For most organizations, the journey to AI compliance will be inexorably linked to the management of third-party AI vendors. A company's reliance on a developer's system does not absolve the deployer of its own responsibilities. The Colorado AI Act's bipartite liability framework effectively extends its extraterritorial reach to any developer who wishes to sell or license a high-risk AI system to a business operating in Colorado. This creates a powerful commercial incentive for deployers to demand compliance from their vendors.

CEOs and CIOs must instruct their procurement and legal teams to embed new, AI-specific requirements into vendor contracts and due diligence processes. Contracts with developers must require them to provide the legally mandated "packet" of documentation, including a general statement of the system's foreseeable uses, information on its training data, known limitations, and risk mitigation measures.

Furthermore, contracts should include audit rights and a contractual obligation for the developer to provide timely disclosures to the Attorney General and all deployers if algorithmic discrimination is discovered.

This proactive approach minimizes legal and operational risks while fostering stronger partnerships built on transparency and accountability. By clearly defining obligations and expectations in contracts, deployers can better ensure the integrity and fairness of the AI systems they utilize. Ultimately, vendors who embrace these compliance standards can differentiate themselves as trusted leaders in a rapidly evolving AI marketplace. This approach shifts the burden of documentation to the developer but becomes a competitive advantage for those vendors who are proactively compliant.

The following table provides an at-a-glance summary of the distinct responsibilities for developers and deployers.

| Obligation | Developer | Deployer |
|---|---|---|
| **Duty of Care** | Yes: Must use reasonable care to protect consumers from known or foreseeable risks of algorithmic discrimination. | Yes: Must use reasonable care to protect consumers from known or foreseeable risks of algorithmic discrimination. |
| **Documentation** | Yes: Must provide documentation to deployers, including purpose, training data, limitations & eval methods. | Yes: Must maintain records of impact assessments for at least three years. |
| **Public Disclosures** | Yes: Must maintain a public website statement summarizing high-risk AI systems and risk management practices. | Yes: Must maintain a public website statement summarizing deployed high-risk AI systems, risk management, and data use. |
| **Impact Assessments** | Not Required. | Yes: Must complete an initial impact assessment and repeat it annually and after substantial modifications. |
| **Consumer Notifications** | Yes: For any AI system interacting with a consumer, if the interaction isn't obvious. | Yes: Must notify consumers before a high-risk system makes a consequential decision and provide detailed disclosure for adverse decisions. |
| **Reporting to AG** | Yes: Must report known or reasonably foreseeable algorithmic discrimination within 90 days. | Yes: Must notify the AG within 90 days of discovering algorithmic discrimination. |

# GOVERNANCE AND ACCOUNTABILITY


This chapter talks about building the internal AI framework.

## Defining and Assigning AI Governance Roles

The Colorado AI Act's requirement for a risk management program that defines "principles, processes, and personnel" necessitates the formalization of AI governance within an organization. This is not a task that can be relegated to a single department. A cross-functional team, led by a designated leader such as a Chief AI Officer or a head of AI Governance, is essential. This team should include representatives from legal, IT, data science, human resources, and compliance to ensure a holistic approach. Clear roles and responsibilities must be defined for each stage of the AI lifecycle, from development and procurement to deployment and monitoring.

## The Board's Fiduciary Duty and Oversight Role

Given the law's potential for significant financial penalties and brand damage, AI governance is no longer a niche technical issue; it is a board-level concern. The board of directors has a fiduciary duty to ensure that the organization has a robust framework in place to mitigate the risks associated with AI, particularly algorithmic discrimination. This oversight role includes approving formal AI governance policies, regularly reviewing the outcomes of impact assessments, and ensuring that adequate resources are allocated to compliance efforts.

## Best Practices for a Centralized AI Governance Program

A fragmented, state-by-state compliance strategy is operationally inefficient and highly risky. The best practice is to design a centralized, harmonized AI governance program that meets or exceeds the Colorado standard and can be scaled for national and international use. The core components of this program should include:

1. **AI System Inventory:** A living document that tracks all AI systems, their purpose, and their classification (e.g., high-risk, low-risk).
2. **Risk-Based Classification:** A clear policy for classifying systems based on their potential for causing algorithmic discrimination.
3. **Impact Assessment Policy:** A formal policy detailing when and how impact assessments are conducted, as well as the template for documentation.
4. **Internal & External Communication Protocols:** Defined processes for internal reporting of issues and for making external disclosures to consumers and the Attorney General.

## Interdepartmental Collaboration: Legal, IT, AI/Data Science and HR

Effective compliance with the Act requires breaking down traditional departmental silos. Legal expertise is needed to interpret the law's nuances and draft policies. IT and data science teams are responsible for the technical execution, including conducting impact assessments and preparing documentation. Human Resources must ensure that AI systems used in employment-related decisions comply with the law's requirements, which includes understanding the implications of disparate impact. Finally, communications teams must manage the public-facing transparency requirements.

## The Role of Third-Party AI Vendors

Most companies will not develop all their AI systems in-house, making vendor management a critical component of a governance program. Companies must review their contracts with AI vendors to ensure that the developer is contractually obligated to provide the required documentation and disclosures. A deployer relying on a third-party AI system is still responsible for its compliance, and without the proper documentation from the developer, they risk being in violation of the Act.

# ENFORCEMENT AND PENALTIES

The Consequences of Non-Compliance can be severe and include Attorney General-led investigations, injunctions and civil penalties of $20K per violation that can reach $ millions in aggregate.

## The Attorney General's Exclusive Enforcement Authority

The Colorado Attorney General (AG) holds the exclusive authority to enforce the Colorado AI Act. The law does not grant a private right of action, which means individual consumers cannot sue companies directly under this specific statute. This legal framework centralizes enforcement power in the hands of the AG, who is also granted plenary rulemaking authority to implement the Act's requirements. This means companies must closely monitor future rulemaking from the AG's office to understand the practical details of compliance.

## The Absence of a Private Right of Action: A Double-Edged Sword

The lack of a private right of action may seem beneficial to businesses as it prevents a flood of individual or class-action lawsuits directly under the Act. However, this is a double-edged sword. While it limits direct litigation risk under the Act, it centralizes enforcement power with the Attorney General. The AG's office can conduct large-scale investigations and audits, and the classification of a violation as a "deceptive trade practice" gives the office broad authority to pursue cases on behalf of the state's consumers.

The risk, therefore, is not from a single consumer but from a state-level regulator who can bring a single case with penalties totaling millions of dollars, making the risk less frequent but potentially more severe.

## Violations as a "Deceptive Trade Practice"

A violation of the Colorado AI Act is classified as a "deceptive trade practice" under the Colorado Consumer Protection Act. This legal classification provides the Attorney General with a powerful enforcement tool and allows them to apply the full range of remedies available under that statute, including injunctions and civil penalties.

## The Financial Penalties: Up to $20,000 per Violation

The financial consequences of non-compliance are substantial, with civil penalties of up to $20,000 per violation. A "violation" can be interpreted per consumer or per transaction, meaning that for a high-volume AI system, fines could quickly escalate into millions of dollars. For example, an AI-powered hiring tool that makes thousands of consequential decisions annually could face catastrophic penalties if found to be in violation.

# SAFE HARBORS AND DEFENSES

This chapter outlines some suggestions on how to build a fortified position under the Colorado AI Act by leveraging safe harbors, recognized frameworks, cure opportunities and rigorous documentation.

## The Rebuttable Presumption: Earning the Shield of "Reasonable Care"

One of the most valuable provisions of the Colorado AI Act is the "rebuttable presumption" of compliance. The law provides that a developer or deployer is presumed to have used "reasonable care" to avoid algorithmic discrimination if they have complied with all of the Act's substantive obligations, including maintaining required documentation and disclosures. This presumption is a powerful legal defense in an enforcement action. It shifts the burden of proof, requiring the Attorney General to prove that the company did not exercise reasonable care, even in the face of its documented compliance efforts.

## Compliance Frameworks as a Strategic Advantage

The Act explicitly points to the **ISO/IEC 42001** and **NIST AI RMF** as recognized standards that can be used to assert an affirmative defense. Adopting one of these frameworks enterprise-wide is a strategic move that not only satisfies the state law but also provides a harmonized approach that can be leveraged across other jurisdictions. The law's design actively incentivizes companies to adopt a globally recognized AI governance framework, creating a powerful alignment between business strategy and regulatory compliance.

## The "Violation Cure" Provision

The law provides an opportunity for a person to cure a violation and receive a defense in an enforcement action. To qualify for this defense, the entity must be in compliance with a recognized risk management framework and have taken specified measures to discover and correct violations of the Act. This provision reinforces the importance of proactive, continuous monitoring and remediation.

## Documentation as a Core Defense

The detailed documentation requirements for both developers and deployers are not merely a bureaucratic burden; they are the foundation of a legal defense. Maintaining comprehensive records of impact assessments, risk mitigation measures, and data governance is essential to proving that "reasonable care" was exercised. This documentation serves as direct evidence of a company's commitment to the Act's principles and is the primary tool for asserting a rebuttable presumption of compliance.

# THE TRUE COSTS OF NON-COMPLIANCE

This chapter makes the business case for action.

## Risks Beyond the Fine

The financial penalties of up to $20,000 per violation under the Colorado AI Act are substantial, but they represent only a fraction of the total risk of non-compliance. The classification of a violation as a "deceptive trade practice" under the Colorado Consumer Protection Act opens a wider aperture of legal and business exposure, triggering a domino effect of financial, operational, and reputational risks. Executives must understand this full spectrum of potential harm to make a compelling business case for proactive governance.

## Financial Risks: Penalties, Litigation Costs, and Audit Expenses

The most immediate and direct risk of non-compliance is financial. The penalties of up to $20,000 per violation can lead to devastating cumulative fines for companies that deploy high-volume AI systems. Non-compliance can trigger costly investigations by the Attorney General's office and lead to the significant expenses associated with remediation efforts.

## Operational Risks: System Halts, Vendor Audits, and Remediation

The Attorney General has the authority to demand audits, impose injunctions that halt the use of an AI system, or require a complete system redesign. These measures can cause significant operational disruptions, leading to unplanned downtime and diverting critical resources from innovation to remediation. Deployers may be forced to audit their AI vendors to ensure they are providing the necessary documentation. The process of remediating a non-compliant AI system can be a significant and unplanned operational drain.

## The Competitive Disadvantage of Inaction

Companies that choose to delay compliance will fall behind their competitors who are proactively building robust AI governance programs. Proactive compliance is an investment in future growth and innovation.

## Reputational and Brand Risks: Public Mistrust and Loss of Market Share

In the era of social media and viral news cycles, a single incident of algorithmic bias can lead to devastating reputational damage. Public perception is paramount, and a scandal can erode consumer and investor trust overnight, leading to a loss of customer loyalty, a decline in sales, and difficulty in attracting talent and securing partnerships. In short, don't be the next headline on CNN!

# BRING IN EXTERNAL ADVISORS

We recommend engaging an AI GRC advisor and external AI Legal counsel for this journey. Qualitative guidance, the opportunity for structured brainstorming and industry experience from external advisors are essential for a defensible program.

Do not rely solely on governance platforms or dashboards; these can track tasks but cannot replace seasoned judgment, board-ready strategy, or credible defense before regulators.

Hiring an AI GRC advisor turns compliance from a checklist into an enterprise capability - policies, controls & documentation - that withstand audits, investigations and market scrutiny. Hiring an experienced AI Legal counsel provides expert guidance on statutory interpretation, risk mitigation, and privileged advice - ensuring compliance decisions are defensible, board-aligned, and able to withstand regulatory scrutiny.

Leverage the expertise of these external advisors throughout the process to build your guardrails, AI policies and disclosures, review vendor contracts, due diligence, and SLAs so platforms and third parties meet disclosure and monitoring duties, deliver board briefings, interpret statutory scope, preserve privilege during assessments, structure consumer notices and appeals, draft and negotiate contract language, and represent the company in Attorney General inquiries and enforcement actions.

# TRUST360™ FOR THE COLORADO AI ACT

Implement governance as a formal project with a named owner, timeline, milestones, and budget so accountability is clear and execution stays on track.

Use a strong methodology like TRUST360™ to mobilize ownership, assess gaps, diagnose risks, implement guardrails, establish vendor controls, enable teams with training and playbooks, then validate and sustain through stress testing, decommissioning, and continuous oversight.

CEOs and CIOs should expect a phased, board-aligned program that inventories AI use, closes control gaps against ISO 42001, embeds consumer notices and appeals with SLAs, and stands up incident-to-cure, recordkeeping, and monitoring for regulator-grade defensibility.


The outcomes are greater stakeholder confidence, faster assessments and AG responses, fewer deployment delays, and clear evidence of reasonable care, with existing tools integrated into a single operating system for AI risk.

Granite Fort Advisory provides the TRUST360™ Methodology as a guided engagement. You can also request a slide-deck on TRUST360™ by sending an email to Engage@GraniteFort.com

# FUTURE OUTLOOK: A SHIFTING REGULATORY LANDSCAPE

## Anticipated Legislative Revisions and "Scaling Back" Proposals

The five-month delay to June 30, 2026, was passed to provide lawmakers with more time to negotiate changes to the Act. The legislative history reveals a strong push from industry to "substantially pare it down". Future debates will likely focus on refining the definition of "algorithmic discrimination," clarifying exemptions, and potentially shifting enforcement authority. While these revisions could make the law less onerous, the repeated failure to reach a consensus on substantive changes suggests that the Act's core principles are unlikely to be abandoned.

## The Ongoing Tension with Federal Policy: Executive Order 14281

The Colorado AI Act's explicit inclusion of "disparate impact" in its definition of algorithmic discrimination is in direct opposition to the federal policy articulated in Executive Order 14281, "Restoring Equality of Opportunity and Meritocracy". This federal policy directs agencies to "deprioritize enforcement" of disparate impact liability in favor of focusing on intentional discrimination. This creates a "two-track compliance environment" for national companies, where the federal government's enforcement posture is less stringent, while Colorado's is more demanding.

The strategic implication is profound: a company could be in full compliance with federal policy and still be in violation of Colorado's law. This is because state attorneys general and private plaintiffs can still pursue disparate impact claims under existing state discrimination laws.12 Therefore, the only safe and prudent strategy is to build a compliance program that can satisfy the more demanding Colorado standard, which, by extension, will also satisfy the less stringent federal standards.

## The Prospect of Federal Preemption and a National Standard

The conflict between state and federal policy raises the possibility of federal preemption, where a national law would supersede state-level regulations. While Executive Order 14281 instructs the Department of Justice to weigh this possibility, a definitive federal law would be needed to create an even playing field and avoid a state-by-state regulatory morass.

Congress has periodically attempted to create such a baseline through the Booker/Wyden Algorithmic Accountability Act (introduce as S.2892/H.R.5628) which would direct the Federal Trade Commission (FTC) to mandate algorithmic impact assessments and mitigation for high-impact uses like employment, credit, housing, education, and health care, with documentation and reporting to the Commission. If enacted, the AAA or potentially other such federal regulations could harmonize core risk-management and transparency practices now appearing in state laws and offer the most direct path to preemption, but it has not advanced beyond introduction, leaving the near-term landscape to state statutes.

Until a national standard is established, companies must assume that Colorado's requirements will stand and use them as a model for a national program. This ensures readiness regardless of where the next state AI law emerges.

## Navigating a Patchwork of State Regulations

Until a national standard is established, companies must assume that Colorado's requirements will stand. The most strategic response is to design a compliance program that can meet the most stringent state requirements, effectively using Colorado's law as a model for a national program. This ensures readiness regardless of where the next state AI law emerges. The Colorado AI Act is, in effect, setting the de facto national benchmark for forward-looking AI governance.

# NEXT STEPS FOR CEOS AND CIOS: THE WAY FORWARD

☑ Elevate AI governance to a boardroom priority by treating AI as a regulated asset.

☑ Work towards establishing a formal AI governance program aligned with global standards such as ISO/IEC 42001.

☑ If you already have an AI governance program, conduct a comprehensive assessment using TRUST360™ or a comparable framework against the Colorado AI Act to identify compliance gaps, risks, and prioritize remediation actions.

**Want to know if your organization is truly ready for Colorado's new AI regulations? Unsure how to fill compliance gaps and reduce AI risks?**

**Contact us** to schedule a TRUST360™ assessment and take confident steps toward compliance and leadership in responsible AI.

**Granite Fort Advisory**
Dallas, TX, United States
Tel: +1-469-713-1511
Engage@GraniteFort.com
www.granitefort.com

GRANITE FORT ADVISORY

**AI Transformation, Governance, Risk & Compliance**

Clarity. Compliance. Confidence.

# APPENDIX 1: RISK **MANAGEMENT POLICY & PROGRAM**

Deployers must implement a risk management policy and program to govern their deployment of a high-risk AI system (Sec. 6-1-1703 (2)). The risk management policy and program must

(1) **specify** the principles, processes, and personnel used to identify and mitigate algorithmic discrimination;

(2) be an iterative process that is **reviewed** and **updated regularly**; and

(3) be **reasonable**, considering factors such as how the framework compares to ISO/IEC 42001 or NIST AI RMF and the size and complexity of the deployer (Sec. 6-1-1703 (2)(a)).

**One risk management policy and program can cover multiple** high-risk AI systems deployed by the deployer (Sec. 6-1-1703 (2)(b)).

Sourced from FPF US Legislation Policy Brief withs minor edits/formatted for clarity:
https://leg.colorado.gov/sites/default/files/images/fpf_legislation_policy_brief_the_colorado_ai_act_final.pdf

# APPENDIX 2: IMPACT ASSESSMENTS

**Annually**, and **within ninety days after a substantial and intentional modification** to a high-risk AI system, a deployer, or a third party contracted to the deployer, must conduct an impact assessment (Sec. 6-1-1703 (3)(a)). As detailed in Sec. 6-1-1703 (3)(b), impact assessments must include, to "the extent reasonably known by or available to the deployer,"

1. **Purpose**: A statement disclosing the system's purpose, intended use cases, deployment context, and benefits (and, if after an intentional and substantial modification, a statement disclosing the extent to which the [AI system] was used in a manner that was consistent with, or varied from, the developer's intended uses);

2. **Risk**: Analysis of whether there are known or reasonably foreseeable risks of algorithmic discrimination and, if so, the nature of those risks and mitigation steps taken;

3. **Data**: A description of categories of data processed as inputs and outputs produced by the system; and an overview of categories of data used to customize the system, if applicable;

4. **Testing**: Metrics used to evaluate the system's performance and known limitations;

5. **Transparency**: A description of transparency measures taken including those to disclose to an individual that the system is in use when it is in use; and

6. **Monitoring**: Description of post-deployment monitoring and user safeguards, such as the deployer's "oversight, use, and learning process" to address issues arising from deployment.

**One impact assessment may cover "a comparable set" of deployed systems**, and an assessment completed for complying with another law or regulation can satisfy the requirements of the CO AI Act if that other assessment "is reasonably similar in scope and effect" to the one required under the Act (Sec. 6-1-1703 (3)(d) & (e)). Impact assessments, and all records concerning each impact assessment, shall be retained for at least three years after the final deployment of the system (Sec. 6-1-1703 (3)(f)).

Sourced from FPF US Legislation Policy Brief with minor edits/formatted for clarity:
https://leg.colorado.gov/sites/default/files/images/fpf_legislation_policy_brief_the_colorado_ai_act_final.pdf

# APPENDIX 3: GLOSSARY OF **KEY TERMS AND ACRONYMS**

**Adversarial Testing**
A method of evaluating AI robustness by introducing crafted inputs designed to produce incorrect or unexpected model outputs.

**Algorithmic Discrimination**
Unjust or prejudicial treatment resulting from AI-driven decisions that disproportionately harm individuals based on protected attributes.

**Bias Mitigation**
Techniques applied during model development—such as data rebalancing or fairness-aware algorithms—to reduce disparate impacts across demographic groups.

**Change Management**
A structured process for controlling modifications to AI systems, including versioning, approvals, and documentation of rationale for retraining or parameter updates.

**Consumer Notification**
A clear, conspicuous disclosure informing users that an AI system is in use, its general purpose, and its limitations before collecting data or making decisions.

**Deployer**
An entity that integrates an AI model into a product, service, or decision-making workflow and is responsible for ongoing risk management and consumer disclosures.

**Developer**
An individual or organization that designs, trains, or modifies an AI system and is accountable for due-care practices and technical documentation.

**Drift Monitoring**
Continuous tracking of model inputs and outputs to detect shifts in data distributions or performance degradation over time.

**Explainability**

Techniques that provide human-understandable insights into how an AI model arrives at its decisions, such as feature-importance scores or counterfactual explanations.

**Governance Committee**

A cross-functional group - often including technology, legal, compliance, and ethics representatives - that oversees AI risk management and policy adherence.

**High-Risk AI System**

Any AI application whose outputs affect legally protected rights or economic interests, such as hiring, lending, insurance, healthcare, housing, or public benefits decisions.

**Impact Assessment**

A documented analysis identifying potential harms, affected populations, likelihood of adverse outcomes, and mitigation plans for a high-risk AI system.

**ISO/IEC 42001**

An international standard specifying requirements for establishing, implementing, maintaining, and continually improving an AI management system.

**Safe Harbor**

Provisions that allow organizations to cure identified compliance deficiencies within a specified time frame before facing enforcement penalties.

**Whistleblower Process**

Established channels and protections that enable employees to report suspected AI-related compliance violations or ethical concerns without fear of retaliation.