eBook

Access

AI Governance and Compliance:

Managing Risk in a Fragmented Regulatory Landscape

AUTHOR

Adam Koonce

ACP, Legal Research Manager, Access

Recently, artificial intelligence has stepped into the spotlight and completely stolen the show. Executives across industries are attracted to the possibilities it offers to streamline operations, improve efficiency, and differentiate their organization. Even in industries previously thought to be static or commoditized, artificial intelligence creates a measurable and lasting impact.

That impact is felt most acutely in areas responsible for managing and governing information. The Records and Information Management (RIM) industry and the organizations it supports are already contending with increasing data volumes and regulatory complexity, and now, the shift in focus towards AI is creating an even more challenging operational reality.

On one hand, AI enables organizations to operate more efficiently. It reduces manual processes, improves accuracy, and allows teams to extract value from information that would otherwise remain underutilized.

On the other hand, these same capabilities introduce new layers of complexity. AI systems process vast amounts of data, often across environments that lack clear visibility or control. They generate new forms of records, including prompts, outputs, logs, and decision artifacts, many of which may fall within existing retention, privacy, and legal discovery obligations.

At the same time, regulatory frameworks are evolving unevenly across jurisdictions. Some governments are implementing comprehensive, risk-based legislation, while others are taking fragmented or sector-specific approaches. Compounding this challenge is the rapid rise of decentralized AI usage as employees are increasingly adopting AI tools outside of approved systems and governance structures.

This eBook examines these dual realities and compounding challenges in depth, while also providing practical guidance for building an AI governance framework that balances innovation with control.

The Emerging Global AI Regulatory Landscape

As artificial intelligence becomes embedded in business operations, governments around the world are working to establish regulatory frameworks that address its risks and implications. However, these efforts vary significantly across jurisdictions, creating a complex and evolving compliance environment.

The European Union: A Risk-Based Approach

The European Union is widely regarded as a leader in AI regulation after adopting the EU Artificial Intelligence Act (EU AI Act) in 2024. The Act establishes a cross-sector, risk-based framework that identifies areas of society where AI systems could pose significant risks to health, safety, or fundamental rights, such as biometric identification, education, employment, critical infrastructure, access to essential services, and the administration of justice.

Importantly, the EU AI Act is designed to remain adaptable over time, allowing these areas of protection to be expanded through commission delegated acts that update the Annex without requiring changes to the core regulation. This flexibility enables the EU to respond quickly to emerging risks and evolving uses of AI across additional domains. From a records and information management perspective, the Act introduces clear documentation and recordkeeping expectations.

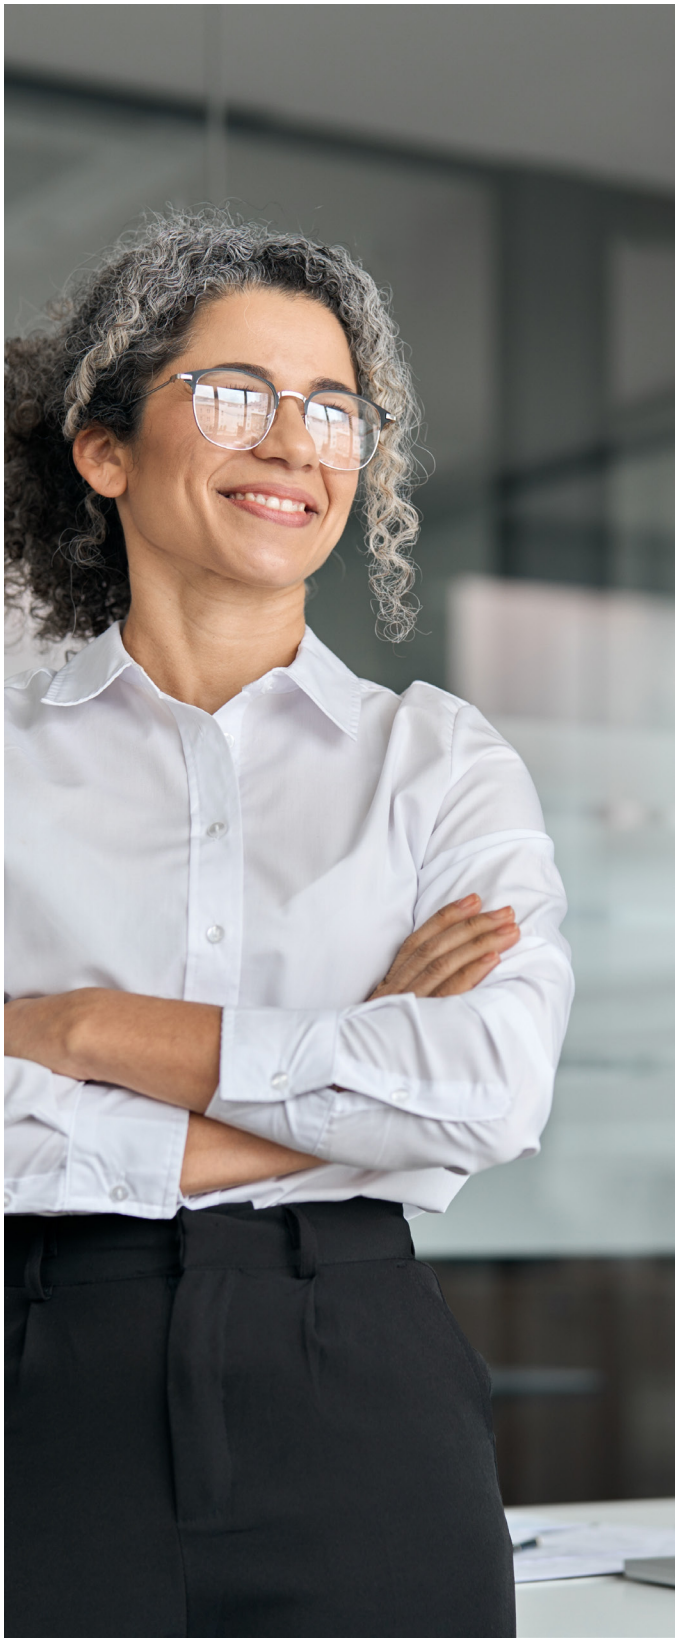

Depending on an organization's role in the lifecycle of an AI product, they may be required to retain:

- **Quality management system documentation**
- **Approved changes to learning models**
- **Technical documentation**
- **EU declarations of conformity**
- **Automatically generated system logs**

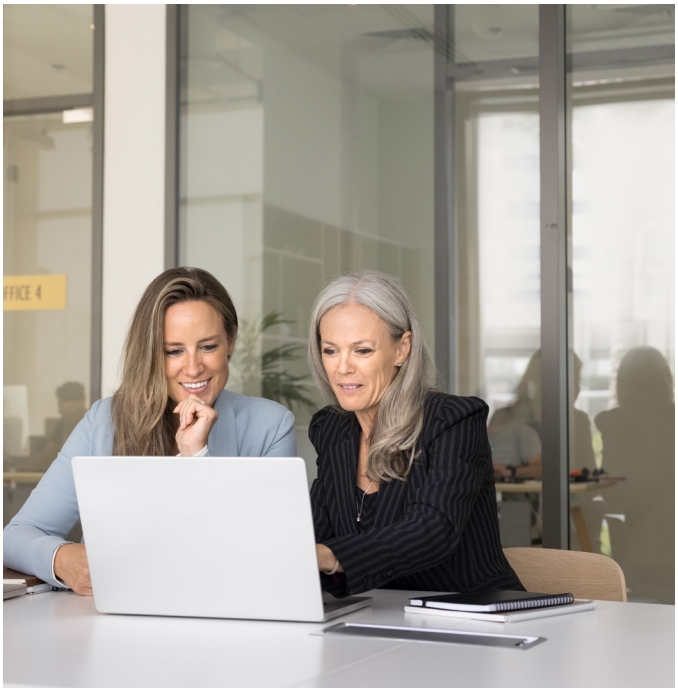

Retention timelines are not uniform across record types. While most documentation must be retained for 10 years, system logs associated with risk events or post-market monitoring have significantly shorter retention periods, typically six months.

Overall, the EU's approach reflects a structured and proactive model for AI governance, one that emphasizes accountability, traceability, and the protection of fundamental rights. This approach is also helping establish a baseline for how AI regulation may develop globally. But while the European Union has taken a centralized and proactive approach, the regulatory landscape in the United States is developing along a very different path.

The United States: Fragmented but Evolving

The United States does not yet have a comprehensive federal AI law. Instead, regulation is emerging through a patchwork of state-level initiatives and sector-specific guidance.

While some states are choosing to focus on business innovation requirements or consumer high-risk protection, others are passing measures to advance technology within their borders. Adding to the complexity, industries like insurance, banking, and healthcare make their own AI guidance, often with no uniformity or alignment.

Colorado has taken the lead with the Colorado Artificial Intelligence Act (CAIA), which took effect in February 2026. The law focuses on high-risk AI systems and requires developers to protect consumers from algorithmic discrimination in areas such as healthcare, housing, insurance, finance, and employment. Mirroring many elements of the EU AI Act, the Colorado provisions for records management around AI are limited to a 3-year requirement to retain impact assessments related to high-risk AI capable of making consequential decisions.

Beyond Colorado, states such as California, Texas, Illinois, and New York have enacted or are advancing AI-specific laws that address algorithmic discrimination, consumer transparency, and governance of high-risk AI systems. Meanwhile, federal lawmakers have introduced hundreds of AI-related bills, though most have not yet resulted in formal legislation. This fragmented regulatory landscape creates compliance uncertainty for organizations attempting to adopt AI responsibly.

Key Differences in Global AI Governance Approaches

For organizations operating across jurisdictions, these differences introduce operational complexity as they're required to align with varying expectations for documentation, accountability, and compliance.

The table below outlines key distinctions between the European Union and United States approaches to AI governance.

An **AI developer** (often called a Provider in regulations) is the entity that creates, designs, or builds the AI system.

An **AI deployer** (often called a User or Operator) is the organization that implements and uses the AI system in real-world operations.

CATEGORY	EUROPEAN UNION	UNITED STATES
Regulatory Structure	Centralized regulation through EU AI Act	Fragmented state-by-state framework
Compliance Timing	Requirements apply before market release	Often evaluated after deployment
Responsibility	Greater focus on AI developers	Greater responsibility on deployers
Risk Classification	Extensive high-risk categories	Narrower definitions in many proposals
Documentation Requirements	Detailed technical documentation and logging	Often lighter unless harm occurs

Regulation is only part of the challenge. Inside many organizations, AI adoption is already happening outside formal oversight structures.

The Rise of Shadow AI

One of the most immediate and difficult governance challenges organizations face today is the emergence of Shadow AI.

Shadow AI refers to the unsanctioned use of artificial intelligence tools outside an organization's approved systems or governance frameworks. These tools may include public chatbots, browser plug-ins, automated writing assistants, or experimental AI platforms used without IT approval.

As employees experiment with AI tools to improve productivity, the risk of exposing sensitive business information grows. For example, a legal team may upload case details for summarization, HR may input employee data to draft evaluations, or engineering may paste proprietary code for debugging. In each case, prompts, responses, uploaded files, and generated outputs may contain confidential information.

Combating Shadow AI with Data Minimization

One of the most effective ways to manage Shadow AI risk is through data minimization. This principle is not new and is already embedded in many global privacy regulations.

For example, GDPR/UK GDPR Article 5(1)(c), says personal data must be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed ('data minimisation')."

States like Virginia, Colorado, and Connecticut have adopted language that closely mirrors GDPR, while others, such as California, Oregon, and Delaware, apply functionally similar rules that focus on proportionality and purpose.

In real life, this looks like sharing only what's necessary. If two paragraphs are enough to generate a summary, there's no reason to upload an entire contract.

Where Shadow AI Risk Appears Across the Enterprise

While data minimization provides a foundational control, its application varies significantly across functions. Each department interacts with AI differently, creating distinct governance risks that must be addressed in context. Understanding how Shadow AI manifests within each role is essential to building controls that are practical, defensible, and aligned with how the business operates.

Legal & Compliance

Legal and compliance teams love AI for clause extraction, case summarization, and issue spotting, but these are precisely the workflows that mix confidential strategy with regulated personal data. For example:

- *Minimization & storage limitation:* If counsel uploads entire case files to a public AI tool, that can violate the GDPR/UK GDPR requirement that data remain "limited to what is necessary" and not kept longer than needed for the specified purpose (Article 5). Prompts and outputs can become records that must sit under your retention and legal hold frameworks, not in personal AI accounts.
- *Health related matters:* If litigation or investigations touch PHI, the HIPAA minimum necessary rule kicks in. That's hard to prove if uploads contain excess identifiers or full medical histories when a redacted excerpt would suffice.

> **What to do:** Make "excerpts only" the default, implement a redaction tool or assistant, and create policy that AI prompts and outputs are records when they capture legal analysis or evidence.

Human Resources

Human resource teams are increasingly required to manage AI across multiple layers of regulation, spanning supranational, national, and sub-national levels. The examples below reflect a small but meaningful portion of government responses to AI in the workplace.

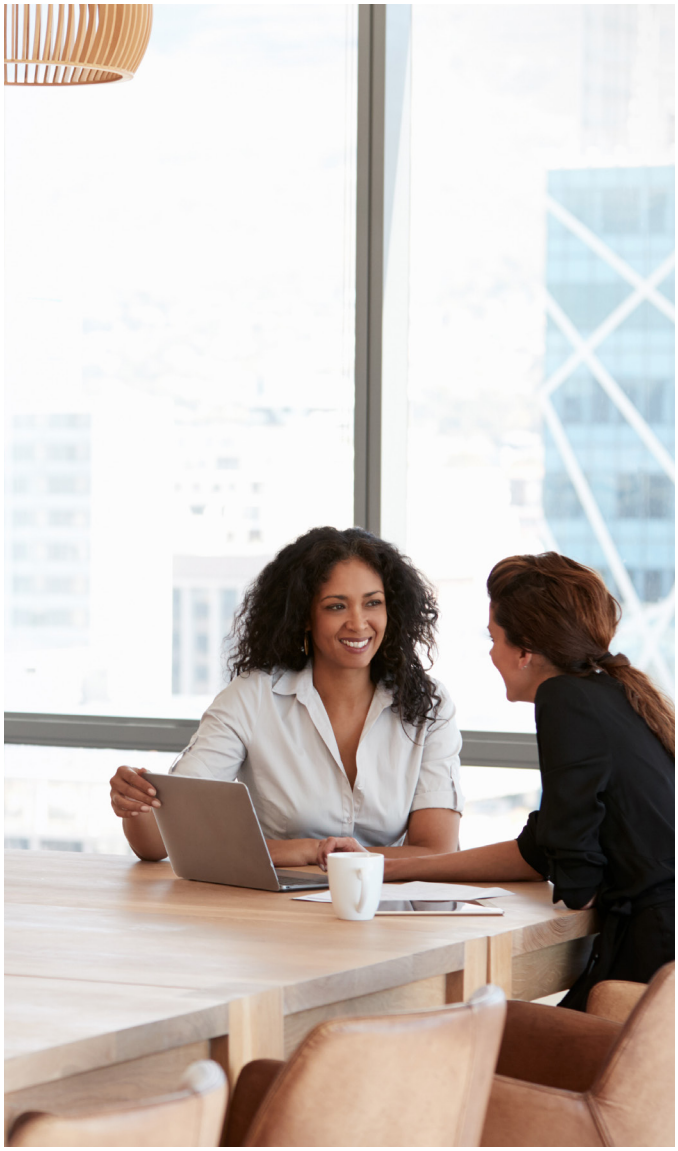

At the supranational level, the EU AI Act explicitly classifies employment-related AI as high risk, including applications used for recruiting, candidate screening, task allocation, performance evaluation, and workforce management. Given the potential impact on workers' fundamental rights, these systems are subject to strict requirements, including documented risk management, logging, transparency, and human oversight.

At the national level, the UK GDPR imposes additional constraints. Articles 22A–D restrict solely automated decisions that produce legal or similarly significant effects, such as hiring or termination. Unless specific legal conditions are met, individuals must be provided with meaningful human review and the ability to contest automated outcomes.

At the sub-national level in the United States, New York City Local Law 144 directly affects employers hiring NYC residents. The law prohibits the use of automated employment decision tools unless a bias audit has been conducted within the past year, the results are publicly available, and advance notice is provided to candidates or employees. Enforcement began in July 2023, and responsibility for compliance remains with the employer, even when a third-party vendor supplies the tool.

What to do: Organizations should limit AI-assisted hiring and performance decisions to approved tools, maintain documented risk assessments and bias audits, ensure human involvement in decision making processes, and confirm that AI outputs and decision rationales are stored in official HR systems of record. This is necessary because a legal action could make this content discoverable and potentially subject to employee data rights oversight.

Engineering and Product Development

Engineering and product development are where trade secrets can be found in source code repositories, proprietary algorithms, design schematics, early-stage prototypes, and product roadmaps that never appear in public filings.

When engineers paste snippets of confidential code into public or unsanctioned AI tools for debugging, optimization, or documentation, they may be disclosing secrets to third-party systems outside the company's governance, audit trail, and contractual safeguards. Similarly, using generative tools to summarize internal roadmaps, analyze unreleased product designs, or refine experimental models can inadvertently expose confidential information.

These seemingly innocuous productivity shortcuts can erode the "reasonable steps" required to maintain trade secret status under the United States' Defend Trade Secrets Act (DTSA) and the European Union's Trade Secrets Directive (2016/943); thus, weakening the organization's ability to claim legal protection if the information leaks, is reused, or surfaces in a competitor's product.

> **What to do:** Use an enterprise AI platform that does not train on your inputs, maintain clear logs of who submitted what information and when, and preserve AI-assisted design decisions in a formal design history file (or equivalent system).

Sales, Marketing, and Customer Success

Sales, marketing, and customer service teams rely on customer data (like order history, payment information, address, phone number, and other personally identifiable information) to operate.

In the EU and UK, teams are increasingly expected to document why specific data elements are needed and how long they should be kept. In the US, although rules vary by state, regulators are similarly questioning whether organizations retain more personal data than necessary for defined operational needs.

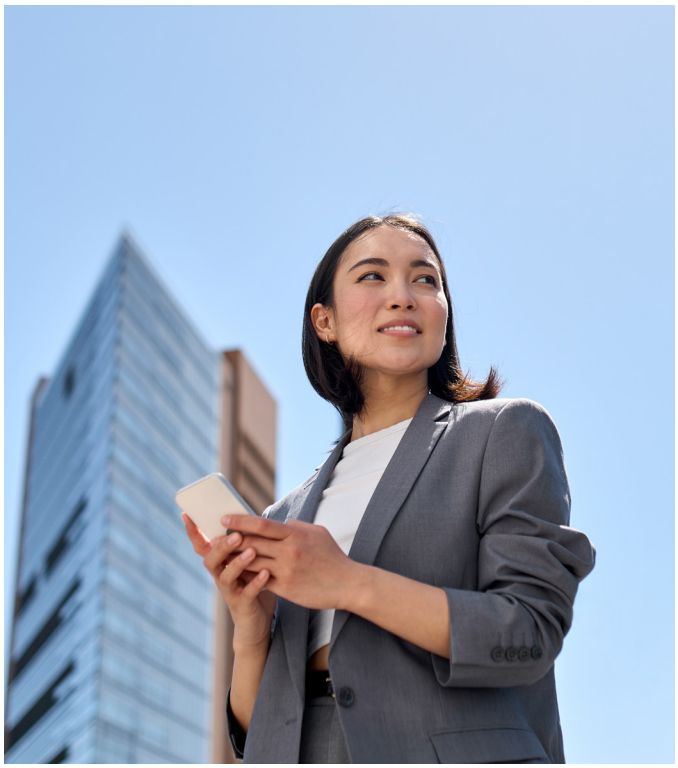

Complicating matters, tools that automatically flag accounts, prioritize renewals, or trigger outreach based on data patterns are increasingly being viewed as automated decision-making, even when people remain involved. In the EU, this can require documenting how decisions are made and ensuring appropriate review. In the US, state laws are beginning to focus on how automated processes affect individuals, especially when they influence access, pricing, or service levels. This means familiar tasks like managing datasets, configuring rules, and supporting workflows are increasingly tied to governance and defensibility, not just efficiency.

> **What to do:** Use only the customer data needed to manage active accounts and renewals, avoid copying records into unsanctioned tools or spreadsheets, and ensure any automated scoring or prioritization can be explained and reviewed if questioned. Limit campaign data to documented purposes, avoid reusing historical or third-party data without clear justification, and ensure targeting or suppression rules are transparent, reviewable, and based on approved datasets. Keep customer records accurate and current, retire outdated usage or sentiment data, and ensure automated alerts or churn signals support human judgment, with decisions documented and defensible.

Building a Practical AI Governance Framework

These function-specific risks highlight a broader reality: AI governance cannot be addressed through policy alone. It requires a structured, organization-wide approach that integrates controls across people, processes, and technology.

Use the following guidance to help you build a practical governance framework for AI adoption within your organization:

Inventory and Document AI Use

- Identify where and how AI tools are being used across departments, including informal or decentralized use (e.g., ChatGPT, internal automation, AI image processing).
- Maintain documentation for each use case, including purpose, data sources, decision logic, and outputs.
- Conduct risk, bias, and impact assessments, particularly for systems that influence decisions about individuals.

2 Establish Policies and Data Controls

- Establish organization-wide guidelines for ethical and compliant AI use, including expectations for handling sensitive data, use of decision-making tools, and mitigation of misinformation risks.
- Implement and enforce data handling controls for AI use, including restrictions on sensitive data, redaction requirements, and the use of approved tools and data sources.
- Define access controls for AI tools and datasets, limiting use based on role, data sensitivity, and approved use cases.
- Maintain an approved inventory of AI tools and implement a vendor review process to evaluate third-party systems for security, privacy, and contractual compliance.

3 Train Teams for AI Awareness

- Provide baseline AI education across all functions, and help non-technical teams understand their legal, regulatory, and ethical obligations.
- Establish ongoing training programs, including annual refreshers and role-specific education, to ensure employees remain aware of evolving risks, policies, and expectations.

4 Align with Standards and Records Management

- Use established standards such as ISO/IEC 42001 and NIST AI RMF as guidance in managing AI-related risks and look to the EU AI Act for best practices, especially if your organization is global.
- Regulators usually evaluate AI systems across lifecycle stages, including design, development, deployment, and ongoing monitoring. Establish governance checkpoints across this lifecycle to assess performance, risk, and compliance before and after deployment.
- Integrate AI-generated content, prompts, and system logs into existing records management frameworks, including retention schedules, legal hold processes, and audit requirements.
- Ensure AI outputs and decision logic are documented and can be explained or justified when subject to audit, regulatory review, or legal challenge.

5 Involve Cross-Functional Stakeholders

- Engage stakeholders from legal, compliance, HR, IT, records management, and business units early in the evaluation and deployment of AI tools.
- Define roles and responsibilities to support accountability and consistent oversight.
- Establish a cross-functional AI governance committee responsible for reviewing and approving AI use cases, overseeing high-risk systems, and maintaining alignment across legal, compliance, IT, records management, and business functions.

6 Monitor, Document, and Respond to Risk

- Develop incident response procedures for AI-related risks, including data exposure, biased outcomes, or improper use of automated decision-making tools.
- Implement monitoring and audit processes to detect high-risk or unauthorized AI usage, including Shadow AI activity.
- Continuously assess third-party AI providers for changes in data handling, model behavior, and contractual terms that may impact compliance or risk exposure.
- Establish metrics and reporting mechanisms to evaluate the effectiveness of AI governance controls, including audit findings, incident trends, and compliance outcomes.

7 Promote Ethical and Defensible AI Use

- Maintain detailed records of how key AI-related decisions are made, and consider any potential harm, fairness, bias, and public trust implications.
- Ensure appropriate human oversight in AI-assisted decision-making processes, particularly where outcomes may have legal or material impact.
- Foster a culture of responsible AI use through clear expectations and leadership alignment.

Taken together, these strategies form the foundation of a defensible AI governance model.

How AI Is Transforming Information Management

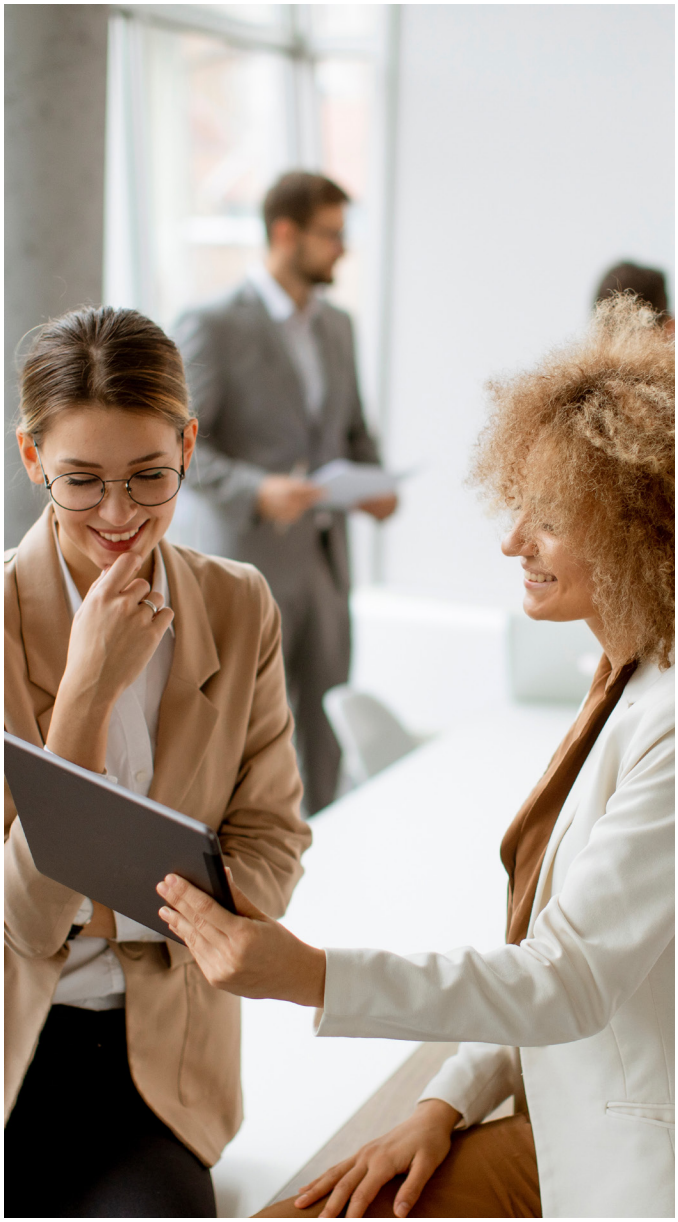

The need for structured AI governance becomes clearer when considering how deeply artificial intelligence is reshaping information management practices. At the operational level, AI streamlines records processes through the automated classification, categorization, and organization of information. This reduces reliance on manual tagging while improving consistency and accuracy across large datasets.

At the access layer, AI enhances search and retrieval through Natural Language Processing (NLP), enabling systems to interpret user intent and surface relevant information from both structured and unstructured sources. This improves information accessibility and supports more efficient decision-making.

From a governance perspective, AI enables more advanced identification of sensitive data, anomaly detection, and compliance monitoring. These capabilities allow organizations to proactively identify and address risks before they escalate into regulatory or legal issues.

AI also expands the analytical value of enterprise data. By identifying patterns, correlations, and anomalies, predictive analytics supports more forward-looking and strategic decision-making.

In addition, AI supports the automation of retention and disposition by aligning records with legal, regulatory, and business requirements. Systems can identify records that have reached retention thresholds and initiate secure disposal, reducing administrative burden while maintaining compliance.

Collectively, these capabilities reinforce the importance of the governance framework outlined above.

As AI becomes more embedded in everyday business operations, it expands not only what organizations can do with their data, but also what they're responsible for managing, documenting, and defending.

However, regulations aren't keeping up; AI is the bullet train being chased by a legislative bicycle. As a result, organizations must navigate growing complexity across compliance, privacy, and operational domains without clear or consistent guidance.

In this environment, governance cannot be reactive. Organizations that fail to establish structured AI governance frameworks risk increased exposure to regulatory scrutiny, data misuse, and operational gaps driven by unmonitored or uncontrolled AI use.

By treating AI-generated data and processes with the same rigor applied to traditional records, organizations can adopt AI with confidence while maintaining control over their information, decisions, and compliance obligations.

For assistance determining which AI related laws and regulations apply to your organization, get in touch with us.

Access

About Access

Access is the largest privately held records and information management services provider worldwide, with operations across the United States, Canada, Central and South America. Access provides transformative services, expertise, and technologies to make organizations more efficient and more compliant. Access helps companies manage and activate their critical business information through offsite storage and information governance services, scanning and digital transformation solutions, document management software, and secure destruction services. Access has been named twelve times to the Inc. 5000, the ranking of fastest-growing private companies in the U.S.